arXiv:1407.7740v3 [cs.DS] 24 Feb 2015

Privacy and Truthful Equilibrium Selection for Aggregative Games

Rachel Cummings* Michael Kearns† Aaron Roth† Zhiwei Steven Wu†

February 26, 2015

Abstract

We study a very general class of games — multi-dimensional aggregative games — which in particular generalize both anonymous games and weighted congestion games. For any such game that is also *large*, we solve the equilibrium selection problem in a strong sense. In particular, we give an efficient *weak mediator*: a mechanism which has only the power to listen to reported types and provide non-binding suggested actions, such that (a) it is an asymptotic Nash equilibrium for every player to truthfully report their type to the mediator, and then follow its suggested action; and (b) that when players do so, they end up coordinating on a particular asymptotic pure strategy Nash equilibrium of the induced complete information game. In fact, truthful reporting is an *ex-post* Nash equilibrium of the mediated game, so our solution applies even in settings of incomplete information, and even when player types are arbitrary or worst-case (i.e. not drawn from a common prior). We achieve this by giving an efficient differentially private algorithm for computing a Nash equilibrium in such games. The rates of convergence to equilibrium in all of our results are inverse polynomial in the number of players n. We also apply our main results to a multi-dimensional market game.

Our results can be viewed as giving, for a rich class of games, a more robust version of the Revelation Principle, in that we work with weaker informational assumptions (no common prior), yet provide a stronger solution concept (ex-post Nash versus Bayes Nash equilibrium). In comparison to previous work, our main conceptual contribution is showing that weak mediators are a game theoretic object that exist in a wide variety of games – previously, they were only known to exist in traffic routing games. We also give the first weak mediator that can implement an equilibrium optimizing a linear objective function, rather than implementing a possibly worst-case Nash equilibrium.

*Computing and Mathematical Sciences, California Institute of Technology; rachelc@caltech.edu. Research performed while the author was visiting the University of Pennsylvania.

†Computer and Information Science, University of Pennsylvania; {mkearns,aaroth,wuzhiwei}@cis.upenn.edu

1 Introduction

Games with a large number of players are almost always played, but only sometimes modeled, in a setting of *incomplete information*. Consider, for example, the problem of selecting stocks for a 401k portfolio among the companies listed in the S&P500. Because stock prices are the result of the aggregate decisions of millions of investors, this is a large multi-player strategic interaction, but it is so decentralized that it is implausible to analyze it in a complete information setting (in which every player knows the types or utilities of all of his opponents), or even in a Bayesian setting (in which every agent shares common knowledge of a prior distribution from which player types are drawn). How players will behave in such interactions is unclear; even under settings of complete information, there remains the potential problem of coordinating or selecting a particular equilibrium among many.

One solution to this problem, recently proposed by Kearns et al. [2014] and Rogers and Roth [2014], is to modify the game by introducing a *weak mediator*, which essentially only has the power to listen and to give advice. Players can ignore the mediator, and play in the original game as they otherwise would have. Alternately, they can use the mediator, in which case they can report their type to it (although they have the freedom to lie). The mediator provides them with a suggested action that they can play in the original game, but they have the freedom to disregard the suggestion, or to use it in some strategic way (not necessarily following it). The goal is to design a mediator such that *good behavior* – that is, deciding to use the mediator, truthfully reporting one's type, and then faithfully following the suggested action – forms an ex-post Nash equilibrium in the mediated game, and that the resulting play forms a Nash equilibrium of the original *complete information* game, induced by the actual (but unknown) player types. A way to approximately achieve this goal – which was shown in Kearns et al. [2014], Rogers and Roth [2014] – is to design a mediator which computes a Nash equilibrium of the game defined by the reported player types under a stability constraint known as *differential privacy* [Dwork et al., 2006]. Prior to our work, this was only known to be possible in the special case of large, unweighted congestion games.

In this paper, we extend this approach to a much more general class of games known as *multi-dimensional aggregative games* (which among other things, generalize both anonymous games and weighted congestion games). In such a game, there is a vector of linear functions of players' joint actions called an *aggregator*. Each player's utility is then a possibly *non-linear* function of the aggregator vector and their own action. For example, in an investing game, the imbalance between buyers and sellers of a stock, which is a linear function of actions, may be used in the utility functions to compute prices, which are a non-linear function of the imbalances (see Section 4 for details). In an anonymous game, the aggregator function represents the number of players playing each action. In a weighted congestion game, the aggregator function represents the total weight of players on each of the facilities. Our results apply to any *large* aggregative game, meaning that any player's unilateral change in action can have at most a bounded influence on the utility of any other player, and the bound on this influence should be a diminishing function in the number of players in the game. Conceptually, our paper is the first to show that *weak mediators* are a game-theoretic object that exists in a large, general class of games: previously, although defined, weak mediators were only known to exist in traffic routing games Rogers and Roth [2014].

This line of work can be viewed as giving robust versions of the Revelation Principle, which can implement Nash equilibria of the complete information game using a "direct revelation mediator," but without needing the existence of a prior type distribution. Compared to the Revelation Principle, which generically requires such a distribution and implements a Bayes Nash equilibrium, truth-telling forms an ex-post Nash equilibrium in our setting. We include a comparison to previous work in Table 1.

Finally, another important contribution of our work is that we are the first to demonstrate the existence of weak mediators (in *any* game) that have the power to optimize over an arbitrary linear function of the actions, and hence able to implement near optimal equilibria under such objective functions, rather than just implementing worst-case Nash equilibria.

Mechanism	**Class of Games**	**Common Prior?**	**Mediator Strength**	**Equilibrium Implemented**
Revelation Principle [Myerson, 1981]	Any Finite Game	Yes	Weak	Bayes Nash
Kearns et al. [2014]	Any Large Game	No	Strong	Correlated
Rogers and Roth [2014]	Large Congestion Games	No	Weak	Nash
This Work	Aggregative Games	No	Weak	Nash

Table 1: Summary of truthful mechanisms for various classes of games and solution concepts. Note that a "weak" mediator does not require the ability to verify player types. A "strong" mediator does. Weak mediators are preferred.

1.1 Our Results and Techniques

Our main result is the existence of a mediator which makes truthful reporting of one's type and faithful following of the suggested action (which we call the "good behavior" strategy) an ex-post Nash equilibrium in the mediated version of any aggregative game, thus implementing a Nash equilibrium of the underlying game of complete information. Unlike the previous work in this line [Kearns et al., 2014, Rogers and Roth, 2014], we do not have to implement an arbitrary (possibly worst-case) Nash equilibrium, but can implement a Nash equilibrium which optimizes any linear objective (in the player's actions) of our choosing. We here state our results under the assumption that any player's action has influence bounded by $(1/n)$ on other's utility, but our results hold more generally, parameterized by the "largeness" of the game.

Theorem 1 (Informal)**.** *In a d-dimensional aggregative game of n players and m actions, there exists a mediator that makes good behavior an η-approximate ex-post Nash equilibrium, and implements a Nash equilibrium of the underlying complete information game that optimizes any linear objective function to within η, where*

$$\eta = O\left(\frac{\sqrt{d}}{n^{1/3}} \cdot \text{polylog}(n, m, d)\right).$$

It is tempting to think that the fact that players only have small influence on one another's utility function is sufficient to make any algorithm that computes a Nash equilibrium of the game a suitable weak mediator, but this is not so (see Kearns et al. [2014] for an example). What we

need out of a mediator is that any single agent's report should have little effect on the *algorithm* computing the Nash equilibrium, rather than on the payoffs of the other players.

The underlying tool that we use is differential privacy, which enforces the stability condition we need on the equilibrium computation algorithm. Our main technical contribution is designing a (jointly) differentially private algorithm for computing approximate Nash equilibria in aggregative games. The algorithm that we design runs in time polynomial in the number of players, but exponential in the dimension of the aggregator function. We note that since aggregative games generalize anonymous games, where the dimension of the aggregator function is the number of actions in the anonymous game, this essentially matches the best known running time for computing Nash equilibria in anonymous games, even non-privately [Daskalakis and Papadimitriou, 2008]. Computing exact Nash equilibria in these games is known to be PPAD complete [Chen et al., 2014].

In the process of proving this result, we develop several techniques which may be of independent interest. First, we give the first algorithm for computing equilibria of multi-dimensional aggregative games (efficient for constant dimensional games) even in the absence of privacy constraints — past work in this area has focused on the single dimensional case Kearns and Mansour [2002], Babichenko [2013]. Second, in order to implement this algorithm privately, we develop the first technique for solving a certain class of linear programs under the constraint of joint differential privacy.

We also give similar results for a class of one-dimensional aggregative games that permit a more general aggregation function and rely on different techniques (Section 5), and we show how our main result can be applied to equilibrium selection in a multi-commodity market (Section 4).

1.2 Related Work

Conceptually, our work is related to the classic Revelation Principle of Myerson [1981], in that we seek to implement equilibrium behavior in a game via a "mediated" direct revelation mechanism. Our work is part of a line, starting with Kearns et al. [2014] and continuing with Rogers and Roth [2014], that attempts to give a more robust reduction, without the need to assume a prior on types. Kearns et al. [2014] showed how to privately compute correlated equilibria (and hence implement this agenda) in arbitrary large games. The private computation of correlated equilibrium turns out to give the desired reduction to a direct revelation mechanism only when the mediator has the power to verify types. Rogers and Roth [2014] rectified this deficiency by privately computing Nash equilibria, but their result is limited to large unweighted congestion games. In this paper, we substantially generalize the class of games in which we can privately compute Nash equilibria (and hence solve the equilibrium selection problem with a direct-revelation mediator).

This line of work is also related to "strategyproofness in the large," introduced by Azevedo and Budish [2012], which has similar goals. In comparison to this work, we do not require that player types be drawn from a distribution over the type-space, do not require any smoothness condition on the set of equilibria of the game, are algorithmically constructive, and do not require our game to be nearly as large. Generally, their results require the number of agents n to be larger than the size of the action set and the size of the type set. In contrast, we only require n to be as large as the *logarithm* of the number of actions, and require no bound at all on the size of the type space (which can even be infinite).

Our work is also related to the literature on mediators in games [Monderer and Tennenholtz, 2003, 2009]. In contrast to our main goal (which is to implement solution concepts of the complete information game in settings of incomplete information), this line of work aims to modify the equilibrium structure of the complete information game. It does so by introducing a mediator, which can coordinate agent actions if they choose to opt in using the mediator. Mediators can

be used to convert Nash equilibria into dominant strategy equilibria [Monderer and Tennenholtz, 2003], or implement equilibrium that are robust to collusion [Monderer and Tennenholtz, 2009]. Ashlagi et al. [2009] considers mediators in games of incomplete information, in which agents can misrepresent their type to the mediators. Our notion of a mediator is related, but our mediators require substantially less power than the ones from this literature. For example, our mechanisms do not need the power to make payments [Monderer and Tennenholtz, 2003], or the power to enforce suggested actions [Monderer and Tennenholtz, 2009]. Like the mediators of Ashlagi et al. [2009], ours are designed to work in settings of incomplete information and so do not need the power to verify agent types — but our mediators are weaker, in that they can only make suggestions (i.e. players do not need to cede control to our weak mediators).

The computation of equilibria in aggregative games (also known as summarization games) was studied in Kearns and Mansour [2002], which gave efficient algorithms and learning dynamics converging to equilibria in the 1-dimensional case. Babichenko [2013] also studies learning dynamics in this class of games and shows that in the 1-dimensional setting, sequential best response dynamics converge quickly to equilibrium. Our paper is the first to give algorithms for equilibrium computation in the multi-dimensional setting, which generalizes many well studied classes of games, including anonymous games. The running time of our algorithm is polynomial in the number of players n and exponential in the dimension of the aggregation function d, which essentially matches the best known running time for equilibrium computation in anonymous games [Daskalakis and Papadimitriou, 2008].

We use a number of tools from differential privacy Dwork et al. [2006], as well as develop some new ones. In particular, we use the advanced composition theorem of Dwork et al. [2010], the exponential mechanism from McSherry and Talwar [2007], and the sparse vector technique introduced by Dwork et al. [2009] (refined in Hardt and Rothblum [2010] and abstracted into its current form in Dwork and Roth [2014]). We introduce a new technique for solving linear programs under joint differential privacy, which extends a line of work (solving linear programs under differential privacy) initiated by Hsu et al. [2014b].

Finally, our work relates to a long line of work initiated by McSherry and Talwar [2007] using differential privacy as a tool and desideratum in mechanism design. In addition to works already cited, this includes Nissim et al. [2012b,a], Xiao [2013], Ghosh and Ligett [2013], Chen et al. [2013], Blum et al. [2014], Kannan et al. [2015] among others. For a survey of this area see Pai and Roth [2013].

2 Model and Preliminaries

2.1 Aggregative Games

Consider an n-player game with action set $\mathcal{A}$ consisting of m actions and a (possibly infinite) type space $\mathcal{T}$ indexing utility functions. Let $\vec{x} = (x_i, \vec{x}_{-i})$ denote a strategy profile in which player i plays action x_i and the remaining players play strategy profile $\vec{x}_{-i}$. Each player i has a utility function, $u\colon \mathcal{T} \times \mathcal{A}^n \to [-1, 1]$, where a player with type t_i experiences utility $u(t_i, \vec{x})$ when players play according to $\vec{x}$. When it is clear from context, we will use shorthand and write $u_i(\vec{x})$ to denote $u(t_i, \vec{x})$, the utility of player i at strategy profile $\vec{x}$.

The utility functions in *aggregative games* can be defined in terms of a multi-dimensional *aggregator* function $S\colon \mathcal{A}^n \to [-W, W]^d$, which represents a compact "sufficient statistic" to compute player utilities. In particular, each player's utility function can be represented as a function only of her own action x_i and the aggregator of the strategy profile $\vec{x}$: $u_i(\vec{x}) = u_i(x_i, S(\vec{x}))$. We also

assume W to be polynomially bounded by n and m. In aggregative games, the function S_k for each coordinate $k \in [d]$, is an additively separable function: $S_k(\vec{x}) = \sum_{i=1}^{n} f_i^k(x_i)$.[1]

Similar to the setting of Kearns and Mansour [2002] and Babichenko [2013], we focus on γ-*aggregative games*, in which each player has a *bounded influence* on the aggregator:

$$\max_i \max_{x_i, x_i' \in \mathcal{A}} \left\| S(x_i, \vec{x}_{-i}) - S(x_i', \vec{x}_{-i}) \right\|_\infty \leq \gamma, \text{ for all } \vec{x}_{-i} \in \mathcal{A}^{n-1}.$$

That is, the greatest change a player can unilaterally cause to the aggregator is bounded by γ. With our motivation to study large games, we assume γ diminishes with the population size n. We also assume that all utility functions are 1-Lipschitz with respect to the aggregator: for all $x_i \in \mathcal{A}$, $|u_i(x_i, s) - u_i(x_i, s')| \leq \|s - s'\|_\infty$.[2]

For γ-aggregative games, we can express the aggregator more explicitly as

$$S_k(\vec{x}) = \gamma \sum_{i=1}^{n} f_i^k(x_i),$$

where $f_i^k(x_i)$ is the influence of player i's action x_i on the k-th aggregator function, and also $|f_i^k(x_i)| \leq 1$ for all actions $i \in [n]$ and $x_i \in \mathcal{A}$. Let $f_{ij}^k = f_i^k(a_j)$, where a_j denotes the j-th action in $\mathcal{A}$.

We say that player i is playing an η-*best response* to $\vec{x}$ if $u_i(\vec{x}) \geq u_i(x_i', \vec{x}_{-i}) - \eta$, for all $x_i' \in \mathcal{A}$. A strategy profile $\vec{x}$ is an η-*pure strategy Nash equilibrium* if all players are playing an η-best response in $\vec{x}$. We also consider *mixed strategies*, which are defined by probability distributions over the action set. For any profile of mixed strategies, given by a product distribution $\vec{p}$, we can define expected utility $u_i(\vec{p}) = \mathbb{E}_{\vec{x} \sim \vec{p}}\, u_i(\vec{x})$ and the expected aggregator

$$S_k(\vec{p}) = \mathop{\mathbb{E}}_{\vec{x} \sim \vec{p}} S_k(\vec{x}) = \gamma \sum_{i=1}^{n} \sum_{j=1}^{m} f_{ij}^k\, p_{ij} = \gamma \left\langle f^k, \vec{p} \right\rangle. \tag{1}$$

The *support* of a mixed strategy p, denoted $\mathrm{Supp}(\vec{p}_i)$, is the set of actions that are played with non-zero probabilities. A mixed strategy profile $\vec{p}$ is a *mixed strategy Nash equilibrium* if $u_i(\vec{p}) \geq \mathbb{E}_{\vec{x}_{-i} \sim \vec{p}_{-i}}\, u_i(x_i', \vec{x}_{-i})$ for all $i \in [n]$ and $x_i' \in \mathcal{A}$.

For each aggregator s, we define the *aggregative best response*[3] for player i to s as $\vec{\mathrm{BA}}_i(s) = \arg\max_{x_i \in A}\{u_i(x_i, s)\}$, breaking ties arbitrarily. We define the η-*aggregative best response set* for player i to s as

$$\eta\text{-}\vec{\mathrm{BA}}_i(s) = \{x_i \in \mathcal{A} | u_i(x_i, s) \geq \max_{x_i'} u_i(x_i', s) - \eta\}$$

to be the set of all actions that are at most η worse than player i's exact aggregative best response.

Remark 1. *Note that best response is played against the other players' actions x_{-i}, but aggregative best response is played against the aggregator value s. Aggregative best response ignores the effect of the player's action on the aggregator, which is bounded by γ; the player reasons about the utility of playing different actions as if the aggregator value were* promised *to be s. Nevertheless, aggregative*

[1] In the economics literature, aggregative games have more restricted aggregator function: $S_k(\vec{x}) = \sum_{i=1}^{n} x_i$. The games we study are more general, and sometimes referred to as *generalized aggregative games*.

[2] Note that the influence that any single player's action has on the utility of others is also bounded by γ. If $\gamma = o(1/n)$, then any player's utility is essentially independent of other players' actions. Therefore, we further assume that $\gamma = \Omega(1/n)$ for the problem to be interesting. This will also simplify some statements.

[3] Sometimes called *best react* [Babichenko, 2013], and *apparent best response* [Kearns and Mansour, 2002].

best response and best response can translate to each other with only an additive loss of γ in the approximation factor. Furthermore, aggregative best responses to different aggregators can translate to each other as long as the corresponding aggregators are close. If $\|s - s'\|_\infty \leq \alpha$, then the actions in η-$\vec{BA}(s)$ are also in $(\eta + 2\alpha)$-$\vec{BA}(s')$. We state these results more formally in the following lemmas.

Lemma 1. *Let $\vec{x}$ be a strategy profile such that player i's action x_i is an η-best response. Then x_i is also an $(\eta + \gamma)$-aggregative best response to $S(\vec{x})$.*

Proof. Let $s = S(\vec{x})$ and $s' = S(a, x_{-i})$ for some deviation $a \neq x_i$. Since x_i is an η-best response, we know that $u_i(x_i, s) \geq u_i(a, s') - \eta$. By the bounded influence of player i, we know that $\|s - s'\|_\infty \leq \gamma$. Also, by the Lipschitz property of u_i, we have that $|u_i(a, s') - u_i(a, s)| \leq \gamma$. It follows that $u_i(a, s') \geq u_i(a, s) - \gamma$, and therefore $u(x_i, s) \geq u_i(a, s) - \gamma - \eta$. □

Lemma 2. *Let $\vec{x}$ be a strategy profile such that every player is playing η-aggregative best response to $S(\vec{x})$. Then we know that each player is playing $(\eta + \gamma)$-best response, and hence $\vec{x}$ forms a $(\eta + \gamma)$-Nash equilibrium.*

Proof. Let $s = S(\vec{x})$ and $s' = S(a, x_{-i})$ for some deviation $a \neq x_i$. Since x_i is an η-aggregative best response, we know $u(x_i, s) \geq u_i(a, s) - \eta$. We know that $\|s - s'\|_\infty \leq \gamma$ by bounded influence of i. Then by the Lipschitz property of u_i, $u_i(a, s) \geq u_i(a, s') - \gamma$. It follows that $u_i(x_i, s) \geq u_i(a, s') - \eta - \gamma$. □

Lemma 3. *Suppose action x_i is an η-aggregative best response to s for player i. Let s' be aggregator such that $\|s - s'\|_\infty \leq \alpha$. Then x_i is an $(\eta + 2\alpha)$-aggregative best response to s'.*

Proof. Let $a \neq x_i$ be some deviation for player i. Since x_i is an η-aggregative best response to s, we have $u_i(x_i, s) \geq u_i(a, s) - \eta$. By the Lipschitz property of u_i, $u_i(x_i, s') \geq u_i(x_i, s) - \alpha$ and also $u_i(a, s) \geq u_i(a, s') - \alpha$. Combining these inequalities, we have $u_i(x_i, s') \geq u_i(a, s') - \eta - 2\alpha$. □

2.2 Mediated Games

We now define games modified by the introduction of a mediator. A mediator is an algorithm $M : (\mathcal{T} \cup \{\perp\})^n \to \mathcal{A}^n$ which takes as input reported types (or $\perp$ for any player who declines to use the mediator), and outputs a suggested action to each player. Given an aggregative game G, we construct a new game G_M induced by the mediator M. Informally, in G_M, players have several options: they can *opt-out* of the mediator (i.e. report $\perp$) and select an action independently of it. Alternately they can *opt-in* and report to it some type (not necessarily their true type), and receive a suggested action r_i. They are free to follow this suggestion or use it in some other way: they play an action $f_i(r_i)$ for some arbitrary function $f_i : \mathcal{A} \to \mathcal{A}$. Formally, the game G_M has an action set $\mathcal{A}_i$ for each player i defined as $\mathcal{A}_i = \mathcal{A}'_i \cup \mathcal{A}''_i$, where

$$\mathcal{A}'_i = \{(t_i, f_i) : t_i \in \mathcal{T}, f_i : \mathcal{A} \to \mathcal{A}\} \quad \text{and } \mathcal{A}''_i = \{(\perp, f_i) : f_i \text{ is constant}\}.$$

Players' utilities in the mediated game are simply their expected utilities induced by the actions they play in the original game. Formally, they have utility functions u'_i: $u'_i(t, f) = \mathbb{E}_{\vec{x} \sim M(t)}[u_i(f(\vec{x}))]$. We are interested in finding mediators such that *good behavior* is an ex-post Nash equilibrium in the mediated game. We first define an ex-post Nash equilibrium.

Definition 1 (Ex-Post Nash Equilibrium). *A collection of strategies* $\{\sigma_i : \mathcal{T} \to \mathcal{A}_i\}_{i=1}^n$ *forms an* η*-approximate ex-post Nash equilibrium if for every type vector* $t \in \mathcal{T}^n$*, and for every player* i *and action* $x_i \in \mathcal{A}_i$*:*

$$u_i'(\sigma_i(t_i), \sigma_{-i}(t_{-i})) \geq u_i'(x_i, \sigma_{-i}(t_{-i})) - \eta$$

That is, it forms an η*-approximate Nash equilibrium for* every *possible vector of types.*

Note that ex-post Nash equilibrium is a very strong solution concept for incomplete information games because it does not require players to know a prior distribution over types.

In a mediated game, we would like players to truthfully report their type, and then faithfully follow the suggested action of the mediator. We call this *good behavior*. Formally, the good behavior strategy is defined as $g_i(t_i) = (t_i, \mathrm{id})$ where $\mathrm{id} : \mathcal{A} \to \mathcal{A}$ is the identity function – i.e. it truthfully reports a player's type to the mediator, and applies the identity function to its suggested action.

In order to achieve this, we use the notion of *joint differential privacy* defined in Kearns et al. [2014] (adapted from *differential privacy*, defined in Dwork et al. [2006] and presented here in Appendix A), as a privacy measure for mechanisms on agents' private data (types). Intuitively, it guarantees that the output to all other agents excluding player i is insensitive to i's private type, so the mechanism protects i's private information from arbitrary coalitions of adversaries.

Definition 2 (Joint Differential Privacy Kearns et al. [2014]). *Two type profiles* t *and* t' *are* i*-neighbors if they differ only in the* i*-th component. An algorithm* $\mathcal{M} : \mathcal{T}^n \to \mathcal{A}^n$ *is* (ε, δ)-joint differentially private *if for every* i*, for every pair of* i*-neighbors* $t, t' \in \mathcal{T}^n$*, and for every subset of outputs* $\mathcal{S} \subseteq \mathcal{A}^{n-1}$*,*

$$\Pr[\mathcal{M}(t)_{-i} \in \mathcal{S}] \leq \exp(\varepsilon) \Pr[\mathcal{M}(t')_{-i} \in \mathcal{S}] + \delta.$$

If $\delta = 0$*, we say that* $\mathcal{M}$ *is* ε-jointly differentially private.

We here quote a theorem of Rogers and Roth [2014], inspired by Kearns et al. [2014] which motivates our study of private equilibrium computation.

Theorem 2 (Rogers and Roth [2014], Kearns et al. [2014]). *Let* M *be a mechanism satisfying* (ε, δ)*-joint differential privacy, that on any input type profile* t *with probability* $1 - \beta$ *computes an* α*-approximate pure strategy Nash equilibrium of the complete information game* $G(t)$ *defined by type profile* t*. Then the "good behavior" strategy* $g = (g_1, \ldots, g_n)$ *forms an* η*-approximate ex-post Nash equilibrium of the mediated game* G_M *for*

$$\eta = \alpha + 2(2\varepsilon + \beta + \delta).$$

Our private equilibrium computation relies on two private algorithmic tools, sparse vector mechanism (called Sparse) and exponential mechanism (called EXP), which allows us to access agents' types in a privacy-preserving manner. (Full details in Appendix A.)

3 Private Equilibrium Computation

Let G be a d-dimensional γ-aggregative game, and $L \colon \mathcal{A}^n \to \mathbb{R}$ be a γ-Lipschitz[4] linear loss function:

$$L(\vec{x}) = \gamma \sum_i \ell_i(x_i) \quad \text{and} \quad L(\vec{p}) = \gamma \mathop{\mathbb{E}}_{\vec{x} \sim \vec{p}} L(\vec{x}) = \gamma \sum_i \langle p_{ij}, \ell_{ij} \rangle.$$

where $0 \leq \ell_i(a_j) \leq 1$ for all actions $a_j \in \mathcal{A}$, and $\ell_{ij} = \ell_i(a_j)$.

[4]This can be achieved by scaling.

Given any $\zeta \geq \gamma\sqrt{8n\log(2mn)}$, let $\mathcal{E}(\zeta)$ be the set of ζ-approximate pure strategy Nash equilibria in the game G,[5] and let

$$\mathrm{OPT}(\zeta) = \min\{L(\vec{x}) \mid \vec{x} \in \mathcal{E}(\zeta)\}.$$

We give the following main result:

Theorem 3. *For any $\zeta \geq \gamma\sqrt{8n\log(2mn)}$, there exists a mediator M that makes good behavior an $(\zeta+\eta)$-approximate ex-post Nash equilibrium of the mediated game G_M, and implements an approximate pure strategy Nash equilibrium $\vec{x}$ of the underlying complete information game with $L(\vec{x}) \leq \mathrm{OPT}(\zeta) + \eta$, where*

$$\eta = O\left(n^{1/3}\gamma^{2/3}\sqrt{d}\cdot \mathrm{polylog}(n,m,d)\right).$$

Recall that the quantity γ is diminishing in n; whenever $\gamma = O(1/n^{1/2+\varepsilon})$ for $\varepsilon > 0$, the approximation factor η tends towards zero as n grows large. Plugging in $\gamma = 1/n$ and $\zeta = \gamma\sqrt{8n\log(2mn)}$ recovers the bound in Theorem 1.

This result follows from instantiating Theorem 2 with an algorithm that computes an approximate equilibrium under joint differential privacy, presented in Algorithm 1 as PRESL (Private Equilibrium Selection).[6] We give here an informal description of our algorithm, absent privacy concerns, and then describe how we implement it privately, deferring the formal treatment to Appendix C.

The main object of interest in our algorithm is the set-valued function

$$\mathcal{V}_\xi(\widehat{s}) = \{S(\vec{p}) \mid \text{for each } i, \mathrm{Supp}(p_i) \subseteq \xi\text{-}\vec{\mathrm{BA}}_i(\widehat{s})\},$$

which maps aggregator values $\widehat{s}$ to the set of aggregator values that arise when players are randomizing between ξ-aggregative best responses to $\widehat{s}$. An approximate equilibrium will yield an aggregator $\widehat{s}$ such that $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$, so we wish to find such a fixed point for $\mathcal{V}_\xi$ (the value of ξ will be determined in the analysis, see Appendix C). Note that *pure strategy* Nash equilibria correspond to such fixed points, but a-priori, it is not clear that fixed points of this function (which may involve mixed strategies) are mixed strategy Nash equilibria. This is because player utility functions need not be linear in the aggregator, and so a best response to the expected value of the aggregator need not be a best response to the corresponding distribution over aggregators. However, as we will show, we can safely round such fixed points to approximate pure strategy Nash equilibria, because the aggregator will be well concentrated under rounding.

For every fixed value $\widehat{s}$, the problem of determining whether $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$ is a linear program (because the aggregator is linear), and although $\mathrm{Supp}(p_i) \subseteq \xi\text{-}\vec{\mathrm{BA}}_i(\widehat{s})$ is not a convex constraint in $\widehat{s}$, the aggregative best responses are fixed for each fixed value of $\widehat{s}$. The first step of our algorithm simply searches through a discretized grid of all possible aggregators $X = \{-W, -W+\alpha, \ldots, W-\alpha\}^d$, and solves this linear program to check if some point $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$. This results in a set of aggregators S that are induced by the approximate equilibria of the game. Let p_{ij} denote the probability that player i plays the j-th action. Then the linear program we need to solve is as

[5] We will show that $\mathcal{E}(\zeta)$ is non-empty for $\zeta \geq \gamma\sqrt{8n\log(2mn)}$ in Appendix C.

[6] We also present the full details of the non-private algorithm to compute equilibrium for aggregative games in Appendix D.

follows:

$$\begin{aligned} &\forall k \in [d], \qquad \widehat{s}_k - \alpha \leq \gamma \sum_{i=1}^{n} \sum_{j=1}^{m} f_{ij}^k p_{ij} \leq \widehat{s}_k + \alpha \\ &\forall i \in [n], \qquad \forall j \in \xi\text{-}\vec{\text{BA}}_i(\widehat{s}), \qquad 0 \leq p_{ij} \leq 1 \\ &\forall i \in [n], \qquad \forall j \notin \xi\text{-}\vec{\text{BA}}_i(\widehat{s}), \qquad p_{ij} = 0 \end{aligned} \tag{2}$$

Next, we need to find a particular equilibrium (an assignment of actions to players) that optimizes our cost-objective function L. This is again a linear program (since the objective function is linear) for each $\widehat{s}$. Hence, for each fixed point $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$ we simply solve this linear program, and out of all of the candidate equilibria, output the one with the lowest cost. Finally, this results in mixed strategies for each of the players, and we round this to a pure strategy Nash equilibrium by sampling from each player's mixed strategy. This does not substantially harm the quality of the equilibrium; because of the low sensitivity of the aggregator, it is well concentrated around its expectation under this rounding. The running time of this algorithm is dominated by the grid search for the aggregator fixed point $\widehat{s}$, which takes time exponential in d. Solving each linear program can be done in time polynomial in all of the game parameters.

Making this algorithm satisfy joint differential privacy is more difficult. There are two main steps. The first is to identify the fixed point $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$ that corresponds the lowest cost equilibrium. There are exponentially in d many candidate aggregators to check, and with naive noise addition we would have to pay for this exponential factor in our accuracy bound. However, we take advantage of the fact that we only need to *output* a single aggregator – the one corresponding to the lowest objective value equilibrium – and so the *sparse vector mechanism* Sparse (described in Appendix A.1) can be brought to bear, allowing us to pay only linearly in d in the accuracy bound.

The second step is more challenging, and requires a new technique: we must actually solve the linear program corresponding to $\widehat{s}$, and output to each player the strategy they should play in equilibrium. The output strategy profile must satisfy joint differential privacy. To do this, we give a general method for solving a class of linear programs (containing in particular, LPs of the form (2)) under joint differential privacy, which may be of independent interest. This algorithm, which we call DistMW (described in Appendix B), is a distributed version of the classic multiplicative weights (MW) technique for solving LPs [Arora et al., 2012]. The algorithm can be analyzed by viewing each agent as controlling the variables corresponding to their own mixed strategies, and performing their multiplicative weights updates in isolation (and ensuring that their mixed strategies always fall within their best response set $\xi\text{-}\vec{\text{BA}}_i(\widehat{s})$). At every round, the algorithm aggregates the current solution maintained by each player, and then identifies a coordinate in which the constraints are far from being satisfied. The algorithm uses the *exponential mechanism* EXP (described in Appendix A.2) to pick such a coordinate while maintaining the privacy of the players' actions. The identification of such a coordinate is sufficient for each player to update their own variables. Privacy then follows by combining the privacy guarantee of the exponential mechanism with a bound on the convergence time of the multiplicative weights update rule. The fact that we can solve this LP in a distributed manner to get joint differential privacy (rather than standard differential privacy) crucially depends on the fact that the sensitivity γ of the aggregator is small. The algorithm DistMW will find a set of strategies that approximately satisfy the linear program – the violation on each coordinate is bounded by

$$E = O\left(\frac{n\gamma^2}{\varepsilon} \operatorname{polylog}\left(n, m, d, \frac{1}{\beta}, \frac{1}{\delta}\right)\right)^{1/2}.$$

The algorithm PRESL has the following guarantee:

ALGORITHM 1: Private Equilibrium Selection via LP: $\mathsf{PRESL}(t, \zeta, L, \varepsilon, \delta, \beta)$

Data: A type vector t, comparator parameter ζ, linear cost function L, privacy parameters (ε, δ), confidence parameter β

Result: An $\widetilde{O}\left(\zeta + \frac{(\sqrt{n\varepsilon}+d)\gamma}{\varepsilon}\right)$-approximate pure strategy Nash equilibrium with cost objective no more than $\mathrm{OPT}(\zeta) + \widetilde{O}\left(\frac{(\sqrt{n\varepsilon}+d)\gamma}{\varepsilon}\right)$

Initialize: discretization resolution $\alpha = E_1 + E_2$, where

$$E_1 = \frac{100\gamma}{\varepsilon}\left((d+1)\log(2W)\log(n) + \log\left(\frac{6}{\beta}\right)\right),$$

$$E_2 = 100\left(\frac{n\gamma^2}{\varepsilon}\log\left(\frac{3d}{\beta}\right)\log(n)\sqrt{\log(m)\ln\left(\frac{1}{\delta}\right)}\right)^{1/2}$$

`/* Find a fixed point` $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$ `that corresponds to the lowest cost equilibrium */`

let $\{a(\widehat{s}, \widehat{y})\} = \mathsf{Sparse}(t, \mathcal{Q}, \alpha + E_1, 1, \varepsilon)$; `//` $\mathcal{Q}$ `formally defined in Appendix C.`

if *all* $a(\widehat{s}, \widehat{y}) = \perp$ **then**

  Abort;

else

  we have $(\widehat{s}, \widehat{y})$ such that $a(\widehat{s}) \neq \perp$;

let $\vec{p} = \mathsf{DistMW}(LP(\widehat{s}, \widehat{y}), \varepsilon, \delta, \alpha, \beta/3)$;

let $\vec{x}$ be an action profile sampled from the product distribution $\vec{p}$;

Output: $\vec{x}$;

Theorem 4. *Let* $\zeta \geq \gamma\sqrt{8n\log(2mn)}, \varepsilon, \delta, \beta \in (0,1)$. $\mathsf{PRESL}(t, \zeta, L, \varepsilon, \delta, \beta)$ *satisfies* $(2\varepsilon, \delta)$*-joint differential privacy, and, with probability at least* $1-\beta$*, computes a* $(\zeta + 12\alpha)$*-approximate pure strategy equilibrium* $\vec{x}$ *such that* $L(\vec{x}) < \mathrm{OPT}(\zeta) + 5\alpha$*, where*

$$\alpha = O\left(\frac{(\sqrt{n\varepsilon}+d)\,\gamma}{\varepsilon}\,\mathrm{polylog}\,(n, m, d, 1/\beta, 1/\delta)\right).$$

We defer the full proof and technical details to Appendix C.

Remark 2. *The running time of this algorithm is exponential in* d*, the dimension of the aggregative game. For games of fixed dimension (where* d *is constant), this yields a polynomial time algorithm. This exponential dependence on the dimension matches the best known running time for (non-privately) computing equilibrium in anonymous games by Daskalakis and Papadimitriou [2008], which is a sub-class of aggregative games.*

Theorem 3 then follows by instantiating Theorem 2 with $\mathsf{PRESL}\left(t, \zeta, L, n^{1/3}\gamma^{2/3}d^{1/2}, \frac{1}{n}, \frac{1}{n}\right)$ – i.e. by setting $\varepsilon = n^{1/3}\gamma^{2/3}d^{1/2}$ and $\delta = \beta = \frac{1}{n}$.

4 An Application to Multi-Commodity Markets

Here we give an application of our main result to a natural market-based game, in which aggregator functions are used to compute non-linear prices.

Consider a market with d types of goods or contracts, which agents can either buy or sell short (i.e. on each contract, an agent can be either long, short, or neutral, and so we can think of actions as being vectors $a \in \mathcal{A} = \{-1, 0, 1\}^d$). In aggregate, the actions of all n players will lead to a

price for each contract, represented by a vector $q \in [0,1]^d$. Agents have (potentially complicated) valuation functions of their positions in the market, modeled as arbitrary functions $v_i \colon \mathcal{A} \to [-d, d]$, and utilities which are quasilinear in money. Note that such valuation functions can model arbitrary complementarity and substitute relationships between contracts. As a result, the equilibria in this market can be complex and diverse, and equilibrium selection becomes a problem.

Central to the game is a *market maker* who sets prices for each contract as a function of the demand. The precise pricing rule that the market maker uses determines the structure of the equilibria of the market. In both real markets and our idealized game, one of the market maker's key objectives is to choose a pricing rule that minimizes his *worst-case loss* — i.e. the loss he might suffer over the buy and sell decisions of the market participants (defined precisely below). A natural and realistic goal is for this loss to be sublinear in the number n of participants or trades. For example, in random-walk models of price movements it is typical for market maker loss to be on the order of $\sqrt{n}$ after n steps or trades. Our model is agnostic to the nature of the commodities being bought and sold — these could, for example, be contracts paying off as a function of the realization of future events, making this a combinatorial prediction market. See Chakraborty and Kearns [2011] and Chen and Vaughan [2010] for analyses of market maker loss in both traditional finance and prediction market models, respectively.

In the following, we show how to phrase the market described here as an aggregative game. We implement a market maker that makes truthful reporting an approximate ex-post Nash equilibrium, and computes an asymptotic Nash equilibrium of the underlying market, all while guaranteeing that the market maker has loss bounded by $O(n^{1/2+\varepsilon})$ per commodity, for any constant $\varepsilon > 0$ (i.e. almost achieving an overall loss $O(d\sqrt{n})$).

4.1 Instantiation of the Market as an Aggregative Game

We will formalize this setting as an n-player γ-aggregative game. The action set of each player is $\mathcal{A} = \{-1, 0, 1\}^d$, where an action is a d-dimensional vector of long, short, or neutral decisions (a portfolio), where 1 and -1 in the k-th coordinate respectively indicate buying and selling a unit of the k-th security. Player i's private type is described by her private valuation function $v_i : \mathcal{A} \to [-d, d]$ that determines her value for any portfolio of d securities, held in positive or negative unit quantities.

Given any strategy profile $\vec{x}$, the *imbalance* in each security is the number agents buying minus the number of agents selling: $I_k(\vec{x}) = \sum_i (x_i)_k$. If the price of security k is q_k and $a_j = 1$, then the player pays q_k; if $a_j = -1$, the player is paid q_k; otherwise, the player receives no payment. The price q_k for each security is a (nonlinear) function of the imbalance vector I parameterized by λ:

$$q_k(I) = \begin{cases} 0 & \text{if } I_k < \frac{-\lambda}{2} \\ I_k/\lambda + 1/2 & \text{if } \frac{-\lambda}{2} \le I_k \le \frac{\lambda}{2} \\ 1 & \text{if } I_k > \frac{\lambda}{2} \end{cases} \tag{3}$$

This simple "hinge" pricing rule is linear with slope λ in a symmetric range of imbalances around 0, and saturates at a price of 1 in the case of overdemand (too many buyers) or 0 in the case of underdemand (too many sellers). We note that all the results discussed here also hold for the standard exponential pricing rule often used in prediction markets Chen and Vaughan [2010].

To apply our main result, we define our aggregator to be $S(\vec{x}) = I(\vec{x})/\lambda$. Conversely, when the aggregator has value s, the imbalance vector is λs and each player i's payoff function (after

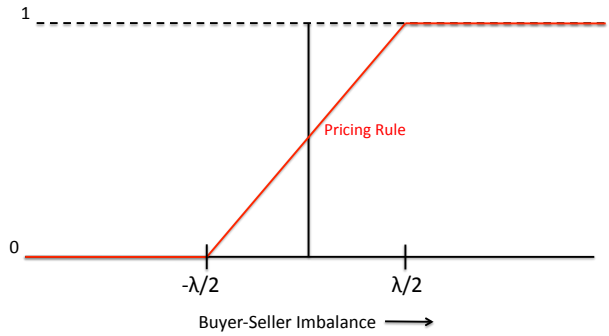

Figure 1: Hinge Pricing Rule

rescaling) for playing action x_i is

$$u_i(x_i, s) = \frac{1}{2d}\left(v_i(x_i) - \langle x_i, q(\lambda s)\rangle\right).^{7}$$

For any fixed action x_i, the payoff u_i is a 1/2-Lipschitz function of price vector q, and the price vector q is an $(1/\lambda)$-Lipschitz function of the imbalance I, which in turn is a λ-Lipschitz function of the aggregator. Therefore, the payoff is a 1-Lipschitz function of the aggregator.[8] Here the range of the aggregator is $[-n/\lambda, n/\lambda]^d$, and each player has bounded influence $\gamma = 1/\lambda$.

4.2 Equilibrium Selection and Market Maker's Loss

For each security k, if I_k is positive, then there are I_k more buyers than sellers, and the market maker must sell to these players at price $q_k(I)$. The market maker will have to pay the maximum price of 1 to procure an extra copy of the item for each player in the worst case, for a potential loss of $\ell_k = I_k(1 - q_k(I))$. Conversely, if I_k is negative, then there are I_k more sellers than buyers, so the market maker must buy from these players at price $q_k(I)$, for a potential loss of $\ell_k = -I_k\, q_k(I)$. In total, the market maker's worst-case loss is $\sum_{k=1}^{d} \ell_k$.

Now consider a mediated game in this market in which the market maker wishes to elicit private valuation functions from all players and make buy/sell recommendations to each player. We have the freedom to set the pricing rule via choice of the parameter λ, but are also faced with a bicriteria problem; we need to set prices to minimize the potential loss of the market maker, while still incentivizing truthful reporting from the players. Here, we demonstrate a trade-off between incentives and market maker's potential loss. First, we have the following lemma which bounds the market maker's loss as a function of λ.

Lemma 4. *The loss ℓ_k for the market maker in each security k under the pricing rule defined in Equation (3) is bounded by $\lambda/16$.*

Proof. Suppose $I_k > 0$. The loss

$$\begin{aligned}\ell_k = I_k(1 - q_k(I)) &\le I_k(1 - I_k/\lambda - 1/2)\\ &= I_k(1/2 - I_k/\lambda)\\ &= -1/\lambda\,(I_k - \lambda/4)^2 + \lambda/16 \le \lambda/16.\end{aligned}$$

[7] We normalize the utility function by $1/2d$ to ensure that $u_i \in [-1, 1]$, so that it fits into the setting defined in Section 2.1.

[8] In fact, the payoff is a 1/2-Lipschitz function of the aggregator, which is a strictly stronger condition.

Suppose $I_k < 0$, we also have

$$\begin{aligned}\ell_k = I_k(-q_k(I)) &\leq I_k(-I_k/\lambda - 1/2) \\ &= -1/\lambda\,(I_k + \lambda/4)^2 + \lambda/16 \leq \lambda/16.\end{aligned}$$

Thus, the loss is always bounded by $\lambda/16$. □

In order to guarantee sub-linear loss in each good, the market maker needs to set $\lambda = o(n)$. Furthermore, since we have a $(1/\lambda)$-aggregative game with each player's action set consisting of 3^d actions, the market maker can use PRESL as a mediator to incentivize truthful reporting.

Corollary 1. PRESL *as our market maker makes each player truthfully reporting their valuation function and following the market maker's recommendation form an η-approximate ex-post Nash equilibrium, where*

$$\eta = O\left(\sqrt{d}\left(\left(\frac{n}{\lambda^2}\right)^{1/3} + \left(\frac{n}{\lambda^2}\right)^{1/2}\right) \cdot \text{polylog}(n, d)\right).$$

This result follows by instantiating Theorem 2 with $\zeta = \sqrt{8nd\log(3n)}/\lambda$. For a fixed number of commodities d, we get asymptotic truthfulness as long as the market maker sets λ to be at least $n^{1/2+\varepsilon}$ for any $\varepsilon > 0$. With this setting of λ, we also guarantee that the market maker experiences worst-case loss at most $O(n^{1/2+\varepsilon})$ per good.

5 Single Dimensional (Quasi)-Aggregative Games

In this section, we consider a more general class of games – *quasi-aggregative games*, in which the aggregator S is not required to have a linear structure as in aggregative games. We focus on γ-quasi-aggregative games with a *one-dimensional* aggregator $S\colon \mathcal{A}^n \to [-W, W]$, and assume the same properties of bounded influence and Lipschitz utilities.[9] We have the following result:

Theorem 5. *Let G be a single dimensional γ-quasi-aggregative game for some $\gamma < 1$. There exists a mediator M that makes good behavior an η-approximate ex-post Nash equilibrium of the mediated game G_M, and implements a Nash equilibrium $\vec{x}$ of the underlying complete information game, where*

$$\eta = O\left(\sqrt{\gamma} \cdot \text{polylog}(n, m)\right).$$

Similar to Section 3, our mediator is a jointly differentially private algorithm that computes an approximate Nash equilibrium. The algorithm is a private implementation of existing algorithms [Kearns and Mansour, 2002, Babichenko, 2013], so we use different techniques for these single-dimensional games. Under certain assumptions, we can also select equilibrium with respect to any Lipschitz objective function of the aggregator (Appendix E.1).

5.1 Private Equilibrium Computation

Our algorithm PSummNash, presented in Algorithm 2, is a privatized version of the SummNash algorithm proposed in Kearns and Mansour [2002], that computes an approximate Nash equilibrium under joint differential privacy.

We will first briefly discuss the main idea of SummNash, with reference to Figure 2. The main object of interest is the function V defined on the aggregator space such that

$$V(s) = S(\vec{\text{BA}}(s)),$$

[9] This is identical to the setting in Kearns and Mansour [2002] and Babichenko [2013].

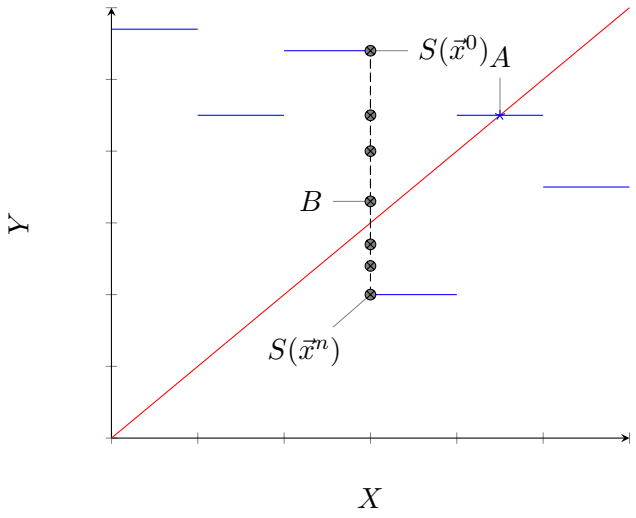

Figure 2: A hypothetical plot of the function V. The ticks on the X axis correspond to the discretized points of aggregators $Z = \{-W, -W + \alpha, \ldots, W - \alpha\}$, and the height of the blue intervals correspond to the values of V evaluated at these points (left endpoint of each interval). The diagonal line correspond to line $Y = X$. The point labeled A is an example of a horizontal crossing the algorithm looks for in stage 1. The column of points including the point labeled B is an example of a vertical crossing the algorithm searches in stage 2. Each point in this column indicates a value of S realized on the sequence of strategy profiles $\vec{x}^0$ through $\vec{x}^n$ defined in (4), while the point B itself is the value of S nearest the diagonal in this walk.

where $\vec{\text{BA}}(s)$ denotes the aggregative best response profile to aggregator s (with each player breaking ties arbitrarily). The algorithm will first discretize the aggregator space $[-W, W]$ into a discrete set $Z = \{-W, -W + \alpha, \ldots, W - \alpha\}$, and evaluate V on each point of Z. In Figure 2, the values of $V(s)$ for each $s \in Z$ corresponds to the horizontal line segments. Then the algorithm finds an equilibrium in two stages.

In the first stage, the algorithm tries to find an approximate fixed point of V in the set Z. Note that if we could identify an aggregator s such that $|V(s) - s| \leq \alpha$, then $\vec{\text{BA}}(s)$ forms an $O(\alpha + \gamma)$-approximate equilibrium by Lemma 2 and Lemma 3. The algorithm simply checks if there exist any $s \in Z$ that satisfy $|V(s) - s| \leq \alpha$. In Figure 2, the existence of an approximate fixed point corresponds to a crossing point A, as the diagonal line corresponds to the relation $Y = X$. If such an s is found, the algorithm will simply suggest that each player i play $\vec{\text{BA}}_i(s)$ and halt.

Otherwise, the algorithm moves to the second stage, where it again iterates over Z, this time to find two adjacent aggregators s and $s + \alpha$ such that $s > V(s) + \alpha$ and $s < V(s + \alpha) - \alpha$. Such pair of $(s, s + \alpha)$ is guaranteed to exist because a failure in the first stage implies that the two endpoints satisfy $V(-W) > -W + \alpha$ and $V(W - \alpha) < W - 2\alpha$. Intuitively, the value of V is "too high" at the lower endpoint, and "too low" at the upper endpoint, so there must some "crossing point" s in the middle where $V(s)$ is close to s.

We can define a sequence of strategy profiles $\mathcal{X} = \{\vec{x}^0, \ldots, \vec{x}^n\}$, where

$$\vec{x}_i^j = \begin{cases} \vec{\text{BA}}_i(s) \text{ if } i \leq j \\ \vec{\text{BA}}_i(s + \alpha) \text{ otherwise} \end{cases} \tag{4}$$

Each profile in $\mathcal{X}$ is a combination of some prefix in $\vec{\text{BA}}(s)$ and $\vec{\text{BA}}(s + \alpha)$. The sequence of aggregators given by the profiles in $\mathcal{X}$ is essentially a walk between $V(s)$ and $V(s + \alpha)$. By the

assumption of bounded influence, changing the action of one player can change the aggregator value by at most γ, so every adjacent step in the walk has length no more than γ. Thus there must be an action profile $\vec{x} \in \mathcal{X}$ such $|V(x) - s| \leq \gamma$. Note that players' actions in x come from both $\vec{\text{BA}}(s)$ and $\vec{\text{BA}}(s+\alpha)$, so all players are playing an $O(\alpha)$-aggregative best response to s in x, and by Lemma 2, x forms an $O(\alpha+\gamma)$-approximate equilibrium.

The primary change we make to SummNash is that our algorithm needs to access the players' aggregative best responses (defined by types) in a privacy-preserving manner. Recall that there are three main parts that require access to the players' private data:

1. for each $s \in Z$, check whether $V(s)$ is close to s;
2. for each $s \in Z$, check whether $s > V(s) + \alpha$ and also $s < V(s+\alpha) - \alpha$; and
3. for each $\vec{x} \in \mathcal{X}$, check whether $S(\vec{x})$ is close to a given $s \in Z$.

To do that, we will first formulate these 3 conditions as 3 different set of queries $\{Q_k\}, \{Q'_k\}$ and $\{Q''_k\}$ (detailed in in Algorithm 2), so that we can check these conditions by checking whether the query values is below some threshold. Note that we only need to identify at most one aggregator or strategy profile that satisfies each condition even though we are answering a large collection queries $(2W/\alpha + n)$. We take advantage of this fact by using Sparse, which gives a good accuracy guarantee, with error scaling *logarithmically* with the number of queries. This error will eventually factor into the approximation factor of output Nash equilibrium.

We here state the formal guarantee of the algorithm, and defer the full proof and technical details to Appendix E.

Theorem 6. PSummNash$(t, \varepsilon, \alpha, \beta)$ *satisfies ε-joint differential privacy, and with probability at least $1-\beta$, computes a $(10\alpha + 2\gamma)$-approximate pure strategy Nash equilibrium as long as*

$$\alpha \geq O\left(\frac{\gamma}{\varepsilon} \operatorname{polylog}(n, m, 1/\beta)\right).$$

Theorem 5 then follows by instantiating Theorem 2 with PSummNash by setting $\varepsilon = \sqrt{\gamma}$, $\beta = 1/n$ and $\alpha = 100\sqrt{\gamma}(\log(12Wn^2))$. (Recall that $\gamma < 1$, so $\sqrt{\gamma}$ dominates γ).

Future Work

The most interesting open question in this line of work is whether there exists a weak mediator that implements good behavior in *every* large game, where we only assume that the influence that any single player's action has on the utility of others is diminishing with the number of players. Recall that in Kearns et al. [2014], it was shown that there exists a *strong* mediator that implements good behavior in any large game, by giving an algorithm that privately computes a correlated equilibrium in any large game. An equivalent result could be shown for weak mediators by giving an algorithm that is able to compute (under the constraint of joint differential privacy) a *Nash* equilibrium, subject only to a largeness condition on the game. Note that such an algorithm would not be expected to be computationally efficient in general. However, at the moment it remains open whether such an algorithm exists at all, independent of efficiency concerns. Finally note that it might be possible to construct weak mediators using tools other than differential privacy – there is no reason why such mediators could not be deterministic. We do not at present have any other similarly general tools for constructing these objects, but results using other tools would be of significant interest.

ALGORITHM 2: $\mathsf{PSummNash}(t, \varepsilon, \alpha, \beta)$

Data: An n-player type vector t, privacy parameter ε, accuracy parameter α, and confidence parameter β

Result: A $(10\alpha + 2\gamma)$-approximate Nash equilibrium

Stage 1

for each $-W/\alpha \leq k \leq W/\alpha - 1$, define a query Q_k on the players' private payoff functions:

$$Q_k = |V(k\alpha) - k\alpha|$$

let $\{a_k\} = \mathsf{Sparse}(t, \{Q_k\}, 4\alpha, 1, \varepsilon/3)$

`// try to find an approximate fixed point for` V

if *we have some* $a_t \neq \perp$ **then**

  Output $\vec{\mathrm{BA}}(t\alpha)$

Stage 2

for each $-W/\alpha + 1 \leq k \leq W/\alpha - 1$, define query Q'_k on the players' private payoff functions:

$$Q'_k = \max\left(\min(0, k\alpha - V((k-1)\alpha)), -2\alpha\right) + \max\left(\min\left(0, V(k\alpha) - k\alpha\right), -3\alpha\right)$$

let $\{a'_k\} = \mathsf{Sparse}(t, \{Q'_k\}, -4\alpha, 1, \varepsilon/3)$

if *all* $a'_k = \perp$ **then**

  Abort.

else

  let l be the index such that $a'_l \neq \perp$

`// define a sequence of strategy profiles for the ‘‘smooth walk’’`

for each $0 \leq j \leq n$, **let** strategy profile x^j be defined as

$$x_i^j = \begin{cases} \vec{\mathrm{BA}}_i(l\alpha) \text{ if } i \leq j \\ \vec{\mathrm{BA}}_i((l-1)\alpha) \text{ otherwise} \end{cases}$$

let query $Q''_j = S(x^j)$

let $\{a''_j\} = \mathsf{Sparse}(t, \{Q''_j\}, \alpha + \gamma/2, 1, \varepsilon/3)$

if *some* $a''_{j'} \neq \perp$ **then**

  Output $x^{j'}$

else

  Abort.

Bibliography

Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.

Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. Mediators in position auctions. *Games and Economic Behavior*, 67(1):2–21, 2009.

Eduardo M Azevedo and Eric Budish. Strategyproofness in the large as a desideratum for market design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, page 55, 2012.

Yaron Azrieli and Eran Shmaya. Lipschitz games. *Math. Oper. Res.*, 38(2):350–357, 2013.

Yakov Babichenko. Best-reply dynamic in large aggregative games. *SSRN abstract 2210080*, 2013.

Avrim Blum, Jamie Morgenstern, Ankit Sharma, and Adam Smith. Privacy-preserving public information for sequential games. *arXiv preprint 1402.4488*, 2014.

Tanmoy Chakraborty and Michael Kearns. Market making and mean reversion. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, pages 307–314, 2011.

Xi Chen, David Durfee, and Anthi Orfanou. On the complexity of nash equilibria in anonymous games. *arXiv preprint arXiv:1412.5681*, 2014.

Yiling Chen and Jennifer Wortman Vaughan. A new understanding of prediction markets via no-regret learning. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10, pages 189–198, 2010.

Yiling Chen, Stephen Chong, Ian A Kash, Tal Moran, and Salil Vadhan. Truthful mechanisms for agents that value privacy. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, EC '13, pages 215–232, 2013.

Constantinos Daskalakis and Christos H. Papadimitriou. Discretized multinomial distributions and Nash equilibria in anonymous games. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, pages 25–34, 2008.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Conference on Theory of Cryptography*, TCC '06, pages 265–284, 2006.

Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil Vadhan. On the complexity of differentially private data release: Efficient algorithms and hardness results. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, STOC '09, pages 381–390, 2009.

Cynthia Dwork, Guy N. Rothblum, and Salil Vadhan. Boosting and differential privacy. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, FOCS '10, pages 51–60, 2010.

Arpita Ghosh and Katrina Ligett. Privacy and coordination: Computing on databases with endogenous participation. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, EC '13, pages 543–560, 2013.

Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, FOCS '10, pages 61–70, 2010.

Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. Private matchings and allocations. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pages 21–30, 2014a.

Justin Hsu, Aaron Roth, Tim Roughgarden, and Jonathan Ullman. Privately solving linear programs. In *Automata, Languages, and Programming*, volume 8572 of *Lecture Notes in Computer Science*, pages 612–624. 2014b.

Ehud Kalai. Large robust games. *Econometrica*, 72, 2004.

Sampath Kannan, Jamie Morgenstern, Aaron Roth, and Zhiwei Steven Wu. Approximately stable, school optimal, and student-truthful many-to-one matchings (via differential privacy). In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1890–1903, 2015.

Michael Kearns and Yishay Mansour. Efficient Nash computation in large population games with bounded influence. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, UAI '02, pages 259–266, 2002.

Michael Kearns, Mallesh Pai, Aaron Roth, and Jonathan Ullman. Mechanism design in large games: Incentives and privacy. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, ITCS '14, pages 403–410, 2014.

Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '07, pages 94–103, 2007.

Dov Monderer and Moshe Tennenholtz. k-Implementation. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, EC '03, pages 19–28, 2003.

Dov Monderer and Moshe Tennenholtz. Strong mediated equilibrium. *Artificial Intelligence*, 173 (1):180–195, 2009.

Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

Kobbi Nissim, Claudio Orlandi, and Rann Smorodinsky. Privacy-aware mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 774–789, 2012a.

Kobbi Nissim, Rann Smorodinsky, and Moshe Tennenholtz. Approximately optimal mechanism design via differential privacy. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, pages 203–213, 2012b.

Mallesh M Pai and Aaron Roth. Privacy and mechanism design. *SIGecom Exchanges*, 12(1):8–29, 2013.

Ryan M. Rogers and Aaron Roth. Asymptotically truthful equilibrium selection in large congestion games. In *Proceedings of the 15th ACM Conference on Economics and Computation*, EC '14, pages 771–782, 2014.

David Xiao. Is privacy compatible with truthfulness? In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, ITCS '13, pages 67–86, 2013.

A Privacy Tools

We first state the formal definition of differential privacy Dwork et al. [2006], which is a measure of the privacy of computations on databases. In our setting, a database $D \in \mathcal{T}^n$ contains n players' private types, which determine their utility functions. Two databases are *neighboring* if they differ only in a single entry.

Definition 3 (Differential Privacy Dwork et al. [2006])**.** *An algorithm $\mathcal{M} : \mathcal{T}^n \to \mathcal{R}$ is* (ε, δ)-differentially private *if for every pair of neighboring databases $D, D' \in \mathcal{T}^n$ and for every subset of possible outputs $\mathcal{S} \subseteq \mathcal{R}$,*

$$\Pr[\mathcal{M}(D) \in \mathcal{S}] \leq \exp(\varepsilon) \Pr[\mathcal{M}(D') \in \mathcal{S}] + \delta.$$

If $\delta = 0$, we say that $\mathcal{M}$ is ε-differentially private.

We will make use of the following composition theorem, which shows how the privacy parameters ε and δ "compose" nicely.

Theorem 7 (Adaptive Composition Dwork et al. [2010])**.** *Let $\mathcal{M} \colon \mathcal{T}^n \to \mathcal{R}^T$ be a T-fold adaptive composition*[10] *of (ε, δ)-differentially private mechanisms. Then $\mathcal{M}$ satisfies $(\varepsilon', T\delta + \delta')$-differential privacy for*

$$\varepsilon' = \varepsilon\sqrt{2T\ln(1/\delta')} + T\varepsilon(e^{\varepsilon} - 1).$$

In particular, for any $\varepsilon \leq 1$, if $\mathcal{M}$ is a T-fold adaptive composition of $(\varepsilon/\sqrt{8T\ln(1/\delta)}, 0)$-differentially private mechanisms, then $\mathcal{M}$ satisfies (ε, δ)-differential privacy.

In the remainder of this section, we review two tools from the differential privacy literature, namely the Sparse Vector Mechanism and the Exponential Mechanism. Both will be used in our algorithms.

A.1 Sparse Vector Mechanism

Our main tool from differential privacy is a slight modification of the sparse vector mechanism from Dwork et al. [2009] (we follow the presentation of Dwork and Roth [2014]). The sparse vector mechanism Sparse takes in a sequence of low-sensitivity queries $\{Q_t\}$ on database D, and a threshold T. The mechanism only outputs answers to those queries with (noisy) answers below the (noisy) threshold,[11] and reports that all other queries were above threshold. There is also an upper bound c on the number of queries that can be answered. If more than c queries have answers below the threshold, the mechanism will abort and not produce an output.

[10] For a more detailed discussion of T-fold adaptive composition, see Dwork et al. [2010].

[11] The Sparse Vector Mechanism as presented in Dwork and Roth [2014] only answers queries with answers above a certain threshold. For the purposes of this paper, we use it instead to answer queries with answers below threshold. This modification does not change the analysis.

This mechanism is especially useful if an analyst is facing a stream of queries and believes that only a small number of the queries will have small answers. The Sparse Vector Mechanism allows the analyst to identify and answer only the "important" queries, without having to incur privacy cost proportional to all queries in the stream.

The *sensitivity* of a query Q, denoted $\Delta(Q)$, is an upper bound over all pairs of neighboring databases on the amount that one entry can affect the answer to the query:

$$\Delta(Q) = \max_{D,D' \text{ s.t. } |D \Delta D'| \leq 1} |Q(D) - Q(D')|$$

Note that because of our restriction to γ-aggregative games in Section 2.1, each coordinate of the aggregative function is γ-sensitive.

ALGORITHM 3: Sparse Vector Mechanism $\mathsf{Sparse}(D, \{Q_t\}, T, c, \varepsilon)$

Data: A private database D, an adaptively chosen stream of queries $\{Q_t\}$ of sensitivity γ, threshold T, total number of numeric answers c, and privacy parameter ε
Result: A stream of answers $\{a_t\}$
  Let $\widehat{T} = T + \text{Lap}\left(\frac{2\gamma}{\varepsilon}\right)$
  let $\sigma = \frac{2c\gamma}{\varepsilon}$
  let count $= 0$
  for *each query* Q_t *on database* D **do**
    Let $\nu_t = \text{Lap}(\sigma)$ and $\widehat{Q}_t = Q_t(D) + \nu_t$
    if $\widehat{Q}_t \leq \widehat{T}$ **then**
      Output $a_t = \widehat{Q}_t$
      Update count $=$ count $+ 1$
      if *count* $\geq c$ **then**
        Abort
    else
      Output $\perp$

Theorem 8 (Dwork et al. [2009])**.** *For any sequence of N queries $Q_1, \ldots, Q_N$ such that $|\{k : Q_k(D) \leq T + \alpha\}| \leq c$,* Sparse *satisfies ε-differential privacy and, with probability at least $1 - \beta$, releases answers such that for all $a_k \in \mathbb{R}$,*

$$|a_k - Q_k(D)| \leq \alpha,$$

and for all $a_k = \perp$,

$$Q_k(D) \geq T - \alpha,$$

where

$$\alpha = \frac{4c\gamma\left(\log N + \log(2c/\beta)\right)}{\varepsilon}.$$

A.2 Exponential Mechanism

The exponential mechanism [McSherry and Talwar, 2007] is a powerful private mechanism for selecting approximately the best outcome from a set of alternatives, where the quality of an outcome is measured by a score function relating each alternative to the underlying data. Let $\mathcal{T}^n$ be the domain of input databases, and $\mathcal{R}$ be the set of possible outcomes, then a score function $q\colon \mathcal{T}^n \times \mathcal{R} \to \mathbb{R}$ maps each database and outcome pair to a real-valued score. The exponential

mechanism EXP instantiated with database D, a score function q, and a privacy parameter ε is defined as

$$\mathsf{EXP}(D, q, \varepsilon) = \text{output } r \text{ with probability proportional to } \exp\left(\frac{\varepsilon q(D, r)}{2\Delta(q)}\right),$$

where $\Delta(q)$ is the global sensitivity of score function q defined as

$$\Delta(q) = \max_{r, D, D' \text{ s.t. } |D \Delta D'| \leq 1} |q(D, r) - q(D', r)|.$$

Then exponential mechanism has the following property:

Theorem 9 (McSherry and Talwar [2007])**.** $\mathsf{EXP}(D, q, \varepsilon)$ *satisfies* ε*-differential privacy and, with probability at least* $1 - \beta$*, outputs an outcome* r *such that*

$$q(D, r) \geq \max_{r'} q(D, r') - \frac{2\Delta(q)\left(\log(|\mathcal{R}|/\beta)\right)}{\varepsilon}.$$

A.3 Billboard Model

In order to prove that our algorithms satisfy joint differential privacy, we rely on a basic but useful framework – the *billboard model*. Algorithms in the billboard model compute some differentially private signal (which can be viewed as being visible on a public billboard); then the output given to each player i is computed as a function only of this private signal, and the private data of agent i. The following lemma shows that algorithms operating in the billboard model satisfy joint differential privacy.

Lemma 5 (Billboard Lemma. Hsu et al. [2014a])**.** *Suppose* $\mathcal{M} \colon \mathcal{T}^n \to \mathcal{R}$ *is* (ε, δ)*-differentially private. Consider any set of functions* $F_i \colon \mathcal{T}_i \times \mathcal{R} \to \mathcal{R}'$*, where* $\mathcal{T}_i$ *is the* i*-th entry of the input data. The composition* $\{F_i(\Pi_i D, \mathcal{M}(D))\}$ *is* (ε, δ)*-jointly differentially private, where* Π_i *is the projection to* i*'s data.*

B Distributed Multiplicative Weights Algorithm

In order to compute an equilibrium privately, we need to solve the linear program in (2) under joint differential privacy. This LP has some nice structural properties which allow this to be possible. In particular, the variables are well partitioned among the n players, such that each player independently controls a set of variables that must form a probability distribution. Each player also has a private restricted feasible set defined by her type (she needs to play an approximate aggregative best response to $\widehat{s}$ according to her private utility function). This motivates us to solve the following more general linear program:

$$\forall k \in [d] \qquad \gamma \sum_{i=1}^{n} \sum_{j=1}^{m} f_{ij}^k \, p_{ij} = \gamma \langle f^k, p \rangle \leq b_k \tag{5}$$

$$\forall i \in [n] \qquad p_i = (p_{i1}, \ldots, p_{im}) \in R_i \subseteq \left\{ x \in \mathbb{R}^m_{\geq 0} \mid \sum_j x_j = 1 \right\} \tag{6}$$

where each $\left|f_{ij}^k\right| \leq 1$. In this LP, there are two types of constraints. Each agent has a private constraint (6) for her own variables, defined by the restricted feasible set R_i. We also have d

cross-agent constraints (5) that require coordination among the agents. Solving this LP under joint differential privacy guarantees that the output variables to the other agents, p_{-i}, is insensitive in agent i's restricted set R_i. Our goal is to find a solution p that approximately satisfies all d cross-agent constraints (5) and exactly satisfies all private constraints (6).

Our algorithm DistMW is essentially a distributed version of the multiplicative weights (MW) update algorithm [Arora et al., 2012].

It proceeds in rounds and has each agent running an instantiation of MW over her own private variables. At each round t, the algorithm collects the variables from all players $p^t = (p_1^t, \ldots, p_n^t) \in \mathbb{R}^{mn}$, and then selects an approximately most violated cross-agent constraint under p^t using the Exponential Mechanism EXP (see Appendix A.2), where the score for a constraint $\gamma\langle f, p^t\rangle \leq b$ is defined as

$$q(p^t, (f, b)) = \gamma\langle f, p^t\rangle - b. \tag{7}$$

Note that each cross-agent constraint takes the same form, so constraint $\gamma\langle f, p^t\rangle \leq b$ can be fully described by the pair (f, b). The mechanism then "broadcasts" the selected constraint (f, b), and each agent i uses the i-th segment of f: $f_i = (f_{i1}, \ldots, f_{im})$ as the loss vector to update her instantiation of the MW distribution. After the re-weighting update at each round, each player projects her vector of variables into her private restricted set R_i, so the solution always satisfies the private constraints. Finally, each agent takes the average of the distributions from all rounds to get her output distribution.

ALGORITHM 4: Distributed Multiplicative Weights for Solving Linear Program: DistMW $(FeasLP, \varepsilon, \delta, \alpha, \beta)$

Data: A feasibility LP $FeasLP$ of with cross-agent constraints of the form (5), private constraints of the form (6), and quality score q of the form (7), privacy parameters (ε, δ), accuracy parameter α, and confidence parameter β

Result: A solution p that satisfies all private constraints and only violates any public constraint by at most α

Initialize p^1 : $p_{ij}^1 = 1/m$ for all $i \in [n]$ and $j \in [m]$

Let $T = \dfrac{16n^2\gamma^2 \log m}{\alpha^2}$ $\qquad \varepsilon_0 = \dfrac{\varepsilon}{2\sqrt{2T\ln(1/\delta)}}$ $\qquad \eta = \alpha/4n\gamma$

For each round $t \in \{1, \ldots, T\}$

Let $(f^t, b^t) = \mathsf{EXP}(p^t, q, \varepsilon_0)$

Each agent i performs MW update: **for** each j

$$\widehat{p}_{ij}^{t+1} = \exp(-\eta \cdot f_{ij}^t) \cdot p_{ij}^t$$

Projection with relative entropy:

$$p_i^{t+1} = \arg\min_{x \in R_i} \mathbf{RE}(x || \widehat{p}_i^{t+1})$$

Output the average vector $\overline{p} = 1/T \sum_{t=1}^{T} p^t$

Theorem 10. DistMW$(\cdot, \varepsilon, \delta, \cdot, \cdot)$ *satisfies* (ε, δ)*-joint differential privacy.*

Proof. Our algorithm works in the *billboard model* introduced by [Hsu et al., 2014a]. In particular, the algorithm posts the violated constraint every round on a *billboard* as a differentially private signal to all agents, such that every agent can see the signal and perform the MW update.

The only sub-routines of DistMW that access the private data (i.e. private constraints) are the constraint selection at every round using the Exponential Mechanism. Thus, our mechanism has T

instantiations of an $(\varepsilon_0, 0)$-differentially private mechanism, where $\varepsilon_0 = \varepsilon/(2\sqrt{2T\ln(1/\delta)})$. By the Adaptive Composition Theorem (Theorem 7 in Appendix A), we know that the selected constraints satisfy (ε, δ)-differential privacy. Note that the i-th component of the output is a function only of the selected constraints and the MW update rule. By Lemma 5, the algorithm satisfies (ε, δ)-joint differential privacy. □

Theorem 11. *Suppose there is a feasible solution to $FeasLP$ with cross-agent constraints (5) and private constraints (6). Then with probability at least $1-\beta$, $\mathsf{DistMW}(FeasLP, \varepsilon, \delta, \alpha, \beta)$ outputs a solution $\overline{p}$ that satisfies all of the private constraints and α-approximately satisfies all of the cross-agent constraints, for*

$$\alpha = O\left(\frac{n\gamma^2}{\varepsilon}\,\mathrm{polylog}\left(n, m, d, \frac{1}{\beta}, \frac{1}{\delta}\right)\right)^{1/2}.$$

Proof. Since each $|f_{ij}^k| \le 1$, the MW algorithm gives a no-regret guarantee for each agent i:

$$\frac{1}{T}\sum_t \langle f_i^t, p_i^t\rangle \le \min_{p_i \in R_i} \frac{1}{T}\sum_t \langle f_i^t, p_i\rangle + \eta + \frac{\log(m)}{T\eta} = \min_{p_i \in R_i} \frac{1}{T}\sum_t \langle f_i^t, p_i\rangle + \frac{\alpha}{2n\gamma}$$

for every agent i. Let $R = R_1 \times \ldots \times R_n$, then the joint play of all n agents satisfy

$$\frac{1}{T}\sum_t \left(\gamma\langle f^t, p^t\rangle - b^t\right) \le \min_{p\in R} \frac{1}{T}\sum_t \left(\gamma\langle f^t, p\rangle - b^t\right) + \alpha/2. \tag{8}$$

Since there is feasible solution to the LP, we know that

$$\min_{p\in R} \frac{1}{T}\sum_t \left(\gamma\langle f^t, p\rangle - b^t\right) \le 0,$$

so from Equation (8),

$$\frac{1}{T}\sum_t \left(\gamma\langle f^t, p^t\rangle - b^t\right) \le \alpha/2. \tag{9}$$

By Theorem 9, with probability at least $1-\beta$, exponential mechanism gives

$$\sum_t \left(\gamma\langle f^t, p^t\rangle - b^t\right) \ge \max_{(f,b)} \sum_t \left(\left(\gamma\langle f, p^t\rangle - b\right) - \frac{2\gamma\log\left(\frac{dT}{\beta}\right)}{\varepsilon_0}\right)$$

and so,

$$\frac{1}{T}\sum_t \left(\gamma\langle f^t, p^t\rangle - b^t\right) \ge \max_{(f,b)} \frac{1}{T}\left[\sum_t \left(\gamma\langle f, p^t\rangle - b\right)\right] - \frac{2\gamma\log\left(\frac{dT}{\beta}\right)}{\varepsilon_0} \tag{10}$$

Combining Equations (9) and (10) with the definition of $\overline{p}$, we get

$$\max_{(f,b)} \left(\gamma\langle f, \overline{p}\rangle - b\right) = \max_{(f,b)} \frac{1}{T}\sum_t \left(\gamma\langle f, p^t\rangle - b\right) \le \frac{2\gamma\log\left(\frac{dT}{\beta}\right)}{\varepsilon_0} + \alpha/2 \le \alpha,$$

as long as $\frac{2\gamma\log\left(\frac{dT}{\beta}\right)}{\varepsilon_0} \le \alpha/2$. Plugging in for parameter ε_0, this condition is equivalent to

$$\alpha^2 \ge \frac{32\sqrt{2}n\gamma^2\log\left(\frac{dT}{\beta}\right)\sqrt{\log m \ln\left(\frac{1}{\delta}\right)}}{\varepsilon}. \tag{11}$$

Plugging in T, we get our desired bound

$$\alpha = \widetilde{O}\left(\frac{n\gamma^2}{\varepsilon}\log\left(\frac{dn}{\beta}\right)\sqrt{\log(m)\ln(1/\delta)}\right)^{1/2}.$$

□

For simplicity, we set the target accuracy to be

$$\alpha = 100\left(\frac{n\gamma^2}{\varepsilon}\log\left(\frac{d}{\beta}\right)\log(n)\sqrt{\log(m)\ln(1/\delta)}\right)^{1/2} \tag{12}$$

when calling $\mathsf{DistMW}(\cdot, \varepsilon, \delta, \alpha, \beta)$.[12]

C Details for PRESL

To recap, our goal is to select an approximate pure strategy equilibrium of a γ-aggregative game such that the objective cost function L is approximately minimized. Before we begin, we need to first define the benchmark OPT for the cost function, that is, the set of equilibria to which we are comparing. We first state the following result showing that approximate equilibrium always exists in large games. It was observed by Kalai [2004] and also proved by Azrieli and Shmaya [2013] using a concentration argument, and we will state the result in terms of aggregative games.

Theorem 12 (Azrieli and Shmaya [2013])**.** *Let G be a γ-aggregative games with n players and m actions, then G admits a $\left(\gamma\sqrt{8n\log(2mn)}\right)$-approximate pure strategy Nash equilibrium.*

Given any $\zeta \geq \gamma\sqrt{8n\log(2mn)}$, let $\mathcal{E}(\zeta)$ be the set of ζ-approximate pure strategy Nash equilibria in the game. We can then define the following benchmark

$$\mathrm{OPT}(\zeta) = \min\{L(\vec{x}) \mid \vec{x} \in \mathcal{E}(\zeta)\}.$$

Now we want to find an $(\zeta + \alpha)$-approximate pure strategy equilibrium that achieves at most $\mathrm{OPT}(\zeta) + O(\alpha)$ using PRESL. The algorithm has two stages. In the first stage, we try to identify an aggregator $\widehat{s} \in X$ such that there exists a mixed strategy profile $\vec{p}$ that satisfies three requirements:

1. all players are randomizing between approximate aggregative best responses to $\widehat{s}$;
2. $S(\vec{p})$ is close to $\widehat{s}$;
3. and $L(\vec{p})$ is close to $\mathrm{OPT}(\zeta)$.

We find such an aggregator by solving linear programs based on the discretized set of candidate aggregators and objective values. In particular, each linear program $LP(\widehat{s}, \widehat{y})$ is defined by a aggregator $\widehat{s} \in X = \{-W, -W+\alpha, \ldots, W-\alpha\}^d$ and a objective value for L: $\widehat{y} \in \{0, \alpha, 2\alpha, \ldots, n\gamma\}$. The sequence of queries we will feed to Sparse is the objective values for these LP's:

$$\mathcal{Q} = \{Q(\widehat{s}, \widehat{y}) \mid \widehat{s} \in X, \widehat{y} \in \{0, \alpha, \ldots, n\gamma\}\},$$

[12] This accuracy level is achievable under mild conditions: as long as $\varepsilon = O(1)$, $\gamma < 1$ and $1/\gamma$ is polynomially smaller than n^{50} (we already assume $\gamma = \Omega(1/n)$), then the α in Equation (12) satisfies Inequality (11).

$$Q(\widehat{s},\widehat{y}) = \min a \tag{13}$$

$$\text{such that} \qquad \forall k, \qquad \gamma \sum_{i=1}^{n}\sum_{j=1}^{m} f_{ij}^{k} p_{ij} \le \widehat{s}_k + a \tag{14}$$

$$\forall k, \qquad -\gamma \sum_{i=1}^{n}\sum_{j=1}^{m} f_{ij}^{k} p_{ij} \le -\widehat{s}_k + a \tag{15}$$

$$L(\vec{p}) \le \widehat{y} + a \tag{16}$$

$$\forall i, \qquad \forall j \in \xi\text{-}\vec{\text{BA}}_i(\widehat{s}), \qquad 0 \le p_{ij} \le 1 \tag{17}$$

$$\forall i, \qquad \forall j \notin \xi\text{-}\vec{\text{BA}}_i(\widehat{s}), \qquad p_{ij} = 0 \tag{18}$$

$$\forall i, \qquad \sum_{j=1}^{m} p_{ij} = 1, \tag{19}$$

where $\xi = \gamma + \zeta + 2\alpha$.

We use Sparse to output the first $(\widehat{s},\widehat{y})$ such that the approximate answer of $Q(\widehat{s},\widehat{y})$ is below the threshold $\alpha + E_1$, where E_1 is the additive error bound for our instantiation of Sparse (given by Theorem 8). This will guarantee the actual value of $Q(\widehat{s},\widehat{y}) \le \alpha + 2E_1$.

During the second stage, the algorithm uses DistMW to compute a mixed strategy profile $\vec{p}$ by solving a modified version of the above LP, denoted $LP(\widehat{s},\widehat{y})$, without the objective (13) and with a replaced by $\alpha + 2E_1$ in the constraints. That is, constraints (14)-(16) are further relaxed as follows:

$$\forall k, \qquad \gamma \sum_{i=1}^{n}\sum_{j=1}^{m} f_{ij}^{k} p_{ij} \le \widehat{s}_k + \alpha + 2E_1 \tag{20}$$

$$\forall k, \qquad -\gamma \sum_{i=1}^{n}\sum_{j=1}^{m} f_{ij}^{k} p_{ij} \le -\widehat{s}_k + \alpha + 2E_1 \tag{21}$$

$$L(\vec{p}) \le \widehat{y} + \alpha + 2E_1. \tag{22}$$

Finally, we will output a pure strategy profile sampled from the distribution defined by $\vec{p}$.

Theorem 4. *Let $\zeta \ge \gamma\sqrt{8n\log(2mn)}, \varepsilon, \delta, \beta \in (0,1)$. PRESL$(t, \zeta, L, \varepsilon, \delta, \beta)$ satisfies $(2\varepsilon,\delta)$-joint differential privacy, and, with probability at least $1-\beta$, computes a $(\zeta + 12\alpha)$-approximate pure strategy equilibrium $\vec{x}$ such that $L(\vec{x}) < \mathrm{OPT}(\zeta) + 5\alpha$, where*

$$\alpha = O\left(\frac{(\sqrt{n\varepsilon} + d)\,\gamma}{\varepsilon}\,\mathrm{polylog}\,(n, m, d, 1/\beta, 1/\delta)\right).$$

Proof. Similar to the privacy proof in Theorem 10, PRESL also works in the Billboard model. Each player's action in the output is a function of only the "broadcast" information and her private type. Note that "broadcast" information comes from both Sparse and DistMW, which together satisfy $(2\varepsilon,\delta)$-differential privacy. By the Billboard Lemma 5, our algorithm satisfies $(2\varepsilon,\delta)$-joint differential privacy.

From Theorem 8, we know that the instantiation Sparse is ε-differentially private, and with probability at least $1-\beta/3$, has additive error bounded by

$$err \le \frac{4\gamma\left((d+1)\log\left(\frac{2W}{\alpha}\right) + \log\left(\frac{6}{\beta}\right)\right)}{\varepsilon} < \frac{100\gamma\left((d+1)\log(2W)\log(n) + \log\left(\frac{6}{\beta}\right)\right)}{\varepsilon} = E_1 < \alpha,$$

as long as $n^{25} > 1/\alpha$. For the rest of the proof, we condition on this level of accuracy, which is the case except with probability at most $\beta/3$. We know there exists an optimal ζ-approximate pure strategy equilibrium $\vec{x}'$ such that $L(\vec{x}') = \text{OPT}(\zeta)$. Note that each player's action x'_i is ζ-best response, and is also ξ-aggregative best response (recall $\xi = \gamma + \zeta + 2\alpha$) to some discretized point $\widehat{s} \in X$, by Lemma 1 and Lemma 3.

Since we set the threshold of Sparse to be $\alpha + E_1$, we are guaranteed that Sparse will output a pair $(\widehat{s}, \widehat{y})$ such that there exists some mixed strategy $\vec{p}$ with

$$\|S(\vec{p}) - \widehat{s}\|_\infty \leq \alpha + 2E_1 \qquad \text{and} \qquad L(\vec{p}) \leq \widehat{y} + \alpha + 2E_1 \leq \text{OPT}(\zeta) + 2\alpha + 2E_1,$$

and where every player only places weight on actions that are a ξ-aggregative best response to $\widehat{s}$.

Then $\vec{p}$ is a feasible solution to the second stage linear program, $LP(\widehat{s}, \widehat{y})$. By Theorem 11, DistMW will output such a solution $\vec{p}$ where

$$\|S(\vec{p}) - \widehat{s}\|_\infty \leq \alpha + 2E_1 + E_2, \text{ and } L(\vec{p}) \leq \text{OPT}(\zeta) + 2\alpha + 2E_1 + E_2,$$

except with probability $\beta/3$, where

$$E_2 = 100\left(\frac{n\gamma^2}{\varepsilon}\log\left(\frac{3d}{\beta}\right)\log(n)\sqrt{\log(m)\ln(1/\delta)}\right)^{1/2} < \alpha.$$

Let $\vec{x}$ be the action profile sampled from the mixed strategy $\vec{p}$. By McDiarmid's inequality, for each coordinate k:

$$\Pr\left[|S_k(\vec{x}) - S_k(\vec{p})| \geq t\right] \leq 2\exp\left(\frac{-2t^2}{n\gamma^2}\right), \text{ and } \Pr\left[|L(\vec{x}) - L(\vec{p})| \geq t\right] \leq 2\exp\left(\frac{-2t^2}{n\gamma^2}\right).$$

The union bound gives

$$\Pr\left[\|S(\vec{x}) - S(\vec{p})\|_\infty \geq t \text{ or } |L(\vec{x}) - L(\vec{p})| \geq t\right] \leq (d+1)2\exp\left(\frac{-2t^2}{n\gamma^2}\right).$$

Then with probability at least $(1 - \beta/3)$, we can guarantee

$$|L(\vec{x}) - L(\vec{p})|, \|S(\vec{x}) - S(\vec{p})\|_\infty \leq \left(\frac{n\gamma^2}{2}\ln\left(\frac{6d+6}{\beta}\right)\right)^{1/2} \triangleq E_3 < \alpha.$$

Overall, we can guarantee the following with probability at least $1 - \beta$

$$\|S(\vec{x}) - \widehat{s}\|_\infty \leq \alpha + 2E_1 + E_2 + E_3 < 4\alpha$$

$$L(\vec{x}) \leq L(\vec{p}) + E_3 \leq \text{OPT}(\zeta) + 2\alpha + 2E_1 + E_2 + E_3 < \text{OPT}(\zeta) + 5\alpha.$$

In $\vec{x}$, all players are playing a ξ-aggregative best response to $\widehat{s}$, so by Lemma 3 they are also playing a $(\xi + 8\alpha)$-aggregative best response to $S(\vec{x})$. By Lemma 2, all players in $\vec{x}$ are playing a $(\xi + \gamma + 8\alpha)$-best response. Since $\gamma < \alpha$ and $\xi + \gamma + 8\alpha < \zeta + 12\alpha$, then $\vec{x}$ is an $(\zeta + O(\alpha))$-pure strategy Nash equilibrium. □

D Non-Private Equilibrium Computation via LP

Here we show that similar techniques to those presented in Appendix C can be used to *non-privately* compute and select approximate equilibrium in γ-aggregative games. In this setting, a better approximation factor is possible because we no longer need to add noise to preserve privacy.

As before, the algorithm NPRESL approximates the aggregator value domain with a discretized grid of all possible aggregators $X = \{-W, -W+\alpha, \ldots, W-\alpha\}^d$, and consider a γ-Lipschitz linear loss function $L\colon \mathcal{A}^n \to \mathbb{R}$:

$$L(\vec{x}) = \gamma \sum_i \ell_i(x_i) \quad \text{and} \quad L(\vec{p}) = \gamma \mathop{\mathbb{E}}_{\vec{x}\sim\vec{p}} L(\vec{x}) = \gamma \sum_i \langle p_{ij}, \ell_{ij} \rangle.$$

where $|\ell_i(a_j)| \le 1$ for all actions $a_j \in \mathcal{A}$, and $\ell_{ij} = \ell_i(a_j)$.

Let $\zeta \ge \gamma\sqrt{8n\log(2mn)}$, and define

$$\mathrm{OPT}(\zeta) = \min\{L(\vec{x}) \mid \vec{x} \in \mathcal{E}(\zeta)\},$$

where $\mathcal{E}(\zeta)$ is the set of ζ-approximate pure strategy equilibria in the game.

We first want to find a mixed strategy profile that could achieve the an objective value no more than $\mathrm{OPT}(\zeta)$. This can be done by solving LP (23) for every $\widehat{s} \in X$, denoted $LP(\widehat{s})$, where $\xi = \zeta + \gamma + 2\alpha$.

$$\begin{aligned}
&\min_p L(p) \\
\forall k \in [d], \qquad & \widehat{s}_k - \alpha \le \gamma \sum_{i=1}^{n}\sum_{j=1}^{m} f_{ij}^k p_{ij} \le \widehat{s}_k + \alpha \\
\forall i \in [n], \qquad & \forall j \in \xi\text{-}\vec{\mathrm{BA}}_i(\widehat{s}), \qquad 0 \le p_{ij} \le 1 \\
\forall i \in [n], \qquad & \forall j \notin \xi\text{-}\vec{\mathrm{BA}}_i(\widehat{s}), \qquad p_{ij} = 0 \\
\forall i \in [n], \qquad & \sum_{j=1}^{m} p_{ij} = 1
\end{aligned} \tag{23}$$

After solving $LP(\widehat{s})$ for all $\widehat{s} \in X$, algorithm NPRESL selects the mixed strategy profile $\vec{p}$ that gives the smallest objective value among all the solutions to the LPs. The algorithm then rounds $\vec{p}$ to get a pure strategy profile $\vec{x}$.

ALGORITHM 5: Non-Private Equilibrium Selection via LP: $\mathsf{NPRESL}(t, \alpha, \beta)$

Data: A type vector t, discretization parameter α, and confidence parameter β
Result: An $\widetilde{O}\left(\alpha + \sqrt{n}\gamma\right)$-approximate pure strategy Nash equilibrium
let s be the aggregator in X that achieves the smallest objective value in $LP(\widehat{s})$;
let $\vec{p}$ be the solution to $LP(s)$;
let $\vec{x}$ be an action profile sampled from the product distribution $\vec{p}$;
Output: $\vec{x}$

Theorem 13. *With probability at least $1-\beta$, NPRESL (t, α, β) computes a $(4\alpha + 2\gamma + 2E)$-approximate pure strategy Nash equilibrium $\vec{x}$ such that $L(\vec{x}) \le \mathrm{OPT} + E$, where*

$$E = O\left(\sqrt{n}\gamma \,\mathrm{polylog}(d, 1/\beta)\right).$$

Proof. Since the (unrounded) mixed strategy profile $\vec{p}$ is a feasible solution to $LP(s)$, we know

$$\|S(\vec{p}) - s\|_\infty \le \alpha.$$

We also know that in $\vec{p}$, every player is playing a ξ-aggregative best response to s, and by Lemma 3, it follows that she is playing a $(\xi + 2\alpha)$-aggregative best response to $S(\vec{p})$.

Let $\vec{x}$ be the realized profile by sampling from $\vec{p}$. By McDiarmid's inequality, for each coordinate k:

$$\Pr\left[|S_k(\vec{x}) - S_k(\vec{p})| \geq t\right] \leq 2\exp\left(\frac{-2t^2}{n\gamma^2}\right), \text{ and } \Pr\left[|L(\vec{x}) - L(\vec{p})| \geq t\right] \leq 2\exp\left(\frac{-2t^2}{n\gamma^2}\right).$$

The union bound gives

$$\Pr\left[\|S(\vec{x}) - S(\vec{p})\|_\infty \geq t \text{ or } |L(\vec{x}) - L(\vec{p})| \geq t\right] \leq (d+1)2\exp\left(\frac{-2t^2}{n\gamma^2}\right).$$

Then with probability at least $1 - \beta$, we can guarantee

$$|L(\vec{x}) - L(\vec{p})|, \|S(\vec{x}) - S(\vec{p})\|_\infty \leq \left(\frac{n\gamma^2}{2}\ln\left(\frac{2d+2}{\beta}\right)\right)^{1/2} \triangleq E.$$

Thus, with probability at least $1 - \beta$, we know that in $\vec{x}$, each player is playing a $(\xi + 2\alpha + 2E)$-aggregative best response to $S(\vec{x})$, and is therefore playing a $(\xi + 2\alpha + 2E + \gamma)$-best response. Hence, we show $\vec{x}$ forms an $(\zeta + 4\alpha + 2\gamma + 2E)$-approximate pure strategy equilibrium.

Note that the optimal ζ-approximate pure strategy equilibrium $\vec{x}'$ with objective value $\mathrm{OPT}(\zeta)$ is also a feasible solution to $LP(\widehat{s})$ for some $\widehat{s} \in X$, so we must have $L(\vec{p}) \leq \mathrm{OPT}(\zeta)$, which implies that $L(\vec{x}) \leq \mathrm{OPT}(\zeta) + E$. □

E Details for Single Dimensional (Quasi)-Aggregative Games

Proof of Theorem 6. Algorithm 2 only accesses the data through three instantiations of Sparse, each of which satisfy $\varepsilon/3$-differential privacy. By the Composition Theorem in Dwork et al. [2006], these three computations compose to satisfy ε-differential privacy. Each player's action in the output strategy profile is a function only of Sparse's output and the player's private data (type). Thus, the algorithm works in the Billboard model, and by the Billboard Lemma 5, it satisfies ε-joint differential privacy.

We now prove that the algorithm computes an approximate Nash equilibrium. Let err be the error in PSummNash due to its calls to Sparse. We know by Theorem 8 that with probability at least $1 - \beta$, all three instantiations of Sparse have error at most

$$err = \frac{100\gamma(\log(2Wn) + \log(6/\beta))}{\varepsilon} \leq \alpha,$$

by our assumption on α. For the rest of the proof we assume this level of accuracy, which is the case except with probability β.

First, consider the case that our algorithm outputs a strategy profile in stage 1. We claim that this gives an $(10\alpha + \gamma)$-approximate Nash equilibrium. Let $\vec{\mathrm{BA}}(k\alpha)$ be the output. Then by the accuracy level of Sparse,

$$|V(k\alpha) - k\alpha| \leq 4\alpha + err \leq 5\alpha.$$

Since each player's action is an aggregative best response to aggregator value $k\alpha$, it is also a 10α-aggregative best response to $V(k\alpha)$ by Lemma 3. Thus, each player is playing a $(10\alpha + \gamma)$-best response as desired by Lemma 2.

Now suppose that the algorithm does not output anything in stage 1. We argue that it will output a $(6\alpha + 2\gamma)$-approximate Nash equilibrium in stage 2.

We first show that the algorithm's second Sparse outputs an index l such that

$$V((l-1)\alpha) < l\alpha < V(l\alpha). \tag{24}$$

Since Sparse failed to output a strategy profile in stage 1, we know that

$$|V(k\alpha) - k\alpha| \geq 4\alpha - err \geq 3\alpha \text{ for all } -W/\alpha \leq k \leq W/\alpha - 1.$$

Since $V(s)$ is in $[-W, W]$ for all s, we know that $V(-W) \geq 3\alpha - W$ and $V(W - \alpha) \leq W - 4\alpha$. Then there must exist an index l such that

$$V((l-1)\alpha) - l\alpha \geq 2\alpha, \text{ and } l\alpha - V(l\alpha) \geq 3\alpha,$$

which implies that $Q'_l = -5\alpha \leq -4\alpha - err$. Such an l satisfies Equation (24) and will also be identified by Sparse as a below threshold query. Note that any index i which does not satisfy Equation (24) must have $Q'_i \geq -3\alpha \geq -4a + err$, so it will not be identified by the Sparse as a below threshold query (since we set the threshold to be -4α).

Finally, we claim that we can find an approximate equilibrium between $\vec{\text{BA}}((l-1)\alpha)$ and $\vec{\text{BA}}(l\alpha)$. Let x^j be the profile defined in Algorithm 2. Then there exists j' such that

$$|S(x^{j'}) - l\alpha| \leq \gamma/2$$

Suppose not. Since $S(x^0) > l\alpha > S(x^n)$, there exists r such that $S(x^r) > l\alpha > S(x^{r+1})$ such that

$$S(x^r) - l\alpha > \gamma/2, \text{ and } l\alpha - S(x^{r+1}) > \gamma/2.$$

However, this violates our bounded influence assumption: $|S(x^{r+1}) - S(x^r)| \leq \gamma$.

Since we set the threshold of the third Sparse to be $\alpha + \gamma/2$, we can find an index j' such that

$$|S(x^{j'}) - l\alpha| \leq 2\alpha + \gamma/2.$$

Note that players in $x^{j'}$ are either playing an aggregative best response to aggregative value $l\alpha$ or $(l-1)\alpha$. It suffices to bound the payoff loss for the latter ones:

$$|S(x^{j'}) - l\alpha| \leq |S(x^{j'}) - (l-1)\alpha| \leq 3\alpha + \gamma/2.$$

Thus an aggregative best response to $(l-1)\alpha$ remains a $(6\alpha + \gamma)$-aggregative best response to $S(x^{j'})$ by Lemma 3. By Lemma 2, each player can only gain at most γ by deviating from an apparent best response, thus each player is playing $(6\alpha + 2\gamma)$-best response, which gives at least a $(10\alpha + 2\gamma)$-approximate equilibrium. □

E.1 Private Equilibrium Computation with a Lipschitz Objective

The PSummNash algorithm presented in Section 5.1 allowed us to compute an approximate Nash equilibrium in any 1-dimensional γ-quasi-aggregative game. However, if the game has multiple approximate equilibria, it does not guarantee the quality of the equilibrium we obtain. In this section we propose an algorithm to select the approximate Nash equilibrium of the highest quality with respect to a given objective. Our algorithm requires the following assumptions on the quasi-aggregative game and objective:

Assumption 1 Each player i has a complete ordering $\succ_i$ over the action set $\mathcal{A}$, where $a \succ_i a'$ if and only if $S(a, x_{-i}) \geq S(a', x_{-i})$ for all $x_{-i} \in \mathcal{A}^{n-1}$. We say player i is playing *optimistically* if she is maximizing the aggregator value with her action, and playing *pessimistically* if minimizing.

Assumption 2 Let $q\colon [-W, W] \to \mathbb{R}$ be a score function that measures the quality of a aggregator value, where $q(s)$ is the quality of aggregator value s. We assume that q is λ-Lipschitz in s.

Kearns and Mansour [2002] shows that any γ-quasi-aggregative game admits a 4γ-approximate pure Nash equilibrium. Given any $\zeta \geq 4\gamma$, let $\mathcal{E}(\zeta)$ be the set of ζ-approximate pure Nash equilibrium. Now define

$$\mathrm{OPT}(\zeta) = \{q(S(\vec{x})) \mid \vec{x} \in \mathcal{E}(\zeta)\}$$

as our benchmark of the equilibrium quality. We will sometimes write OPT for $\mathrm{OPT}(\zeta)$ when the context is clear. We show that our algorithm can compute an $(\zeta + O(\gamma))$-approximate equilibrium with quality at least $\mathrm{OPT} + O(\lambda\gamma)$.

Similar to Algorithm 2, this algorithm also iterates through all values s in the discretized set $Z = \{-W, -W + \alpha, \ldots, W - \alpha\}$, with players submitting their approximate aggregative best response sets to s, ξ-$\vec{\mathrm{BA}}(s)_i$. Let $X(s) = \{\vec{x} \mid$ each player is playing some $x_i \in \xi\text{-}\vec{\mathrm{BA}}(s)_i\}$, where $\xi = 2\alpha + \gamma + \zeta$. We are searching for an approximate pure strategy Nash equilibrium $x' \in X(s)$ such that $|S(x') - s| \leq \alpha$. Note that the cardinality of $X(s)$ can potentially be $\Omega(m^n)$, but by Assumption 1, we can compute the upper and lower bound efficiently

$$S_{max}(s) = \max_{x \in X(s)} S(x), \text{ and } S_{min}(s) = \min_{x \in X(s)} S(x)$$

by asking players to play optimistically and pessimistically, respectively.

If we find such an s that has an approximate Nash equilibrium strategy profile, there are three cases: either $S_{max}(s)$ or $S_{min}(s)$ is close to s, or $s \in [S_{min}(s), S_{max}(s)]$. If we are in the first two cases, the algorithm simply outputs the corresponding to the optimistic or pessimistic strategy profile, respectively. In the third case, the algorithm perform the same smooth walk as in Algorithm 2 from the optimistic profile to the pessimistic profile, and outputs an intermediate profile x' such that $|S(x') - s|$ is small.

Since we are interested in computing the *best* equilibrium, we iterate through aggregators in order of their quality score. Let $s_1 \succ s_2 \succ \ldots \succ s_{2W/\alpha}$ be an ordering over the set of discretized aggregator values Z, such that $q(s_i) \geq q(s_{i+1})$. The algorithm will sequentially consider s_i according to this ordering $\succ$, to compute an approximate equilibrium with aggregator value that maximizes q.

Theorem 14. *For any $\zeta \geq 4\gamma$, Algorithm 6 satisfies ε-joint differential privacy, and with probability at least $1 - \beta$, outputs a $(10\alpha + 3\gamma + \zeta)$-approximate Nash equilibrium with quality score at least $\mathrm{OPT}(\zeta) - 5\alpha\lambda$, for any*

$$\alpha \geq O\left(\frac{\gamma}{\varepsilon} \operatorname{polylog}(n, m, 1/\beta)\right).$$

Proof. Algorithm 6 only accesses the data through four instantiations of Sparse, each of which answers at most one query with $\varepsilon/4$-differential privacy. Again, by the Composition Theorem in Dwork et al. [2006], these privacy parameters compose so that the strategy profile selection as a public message satisfies ε-differential privacy. The action of each player is a function of only the public message and her own private payoff data (type), so by Lemma 5, Algorithm 6 satisfies ε-joint differential privacy.

ALGORITHM 6: Private Equilibrium Selection with a Lipschitz Objective

Data: An n-player type vector t, comparator equilibrium class parameter ζ, privacy parameter ε, accuracy parameter α, and confidence parameter β

Result: $(10\alpha + 3\gamma + \zeta)$-approximate Nash equilibrium with quality score at least $\mathrm{OPT}(\zeta) - 5\alpha\lambda$

Initialize: G such that $q(G) = -W$

for *each aggregator value* $s \in \{s_1, s_2, \ldots, s_{2W/\alpha}\}$ **do**

  let $X(s) = \{\vec{x} \mid \text{each player has } x_i \in (2\alpha + \gamma + \zeta)\text{-}\vec{\mathrm{BA}}(s)_i\}$

  let $S_{max}(s) = \max_{x \in X(s)} S(\vec{x})$ and $S_{min}(s) = \min_{x \in X(s)} S(\vec{x})$

for $1 \le k \le 2W/\alpha$ **do**

  let queries

$$Q_k = |S_{max}(s_k) - s_k| \qquad Q'_k = |S_{min}(s_k) - s_k|$$

$$Q''_k = \max\left(\min\left(S_{min}(s_k) - s_k, 0\right), -2\alpha\right) + \max\left(\min\left(s_k - S_{max}(s_k), 0\right), -2\alpha\right)$$

let $\{a_k\} = \mathsf{Sparse}(t, \{Q_k\}, 3\alpha, 1, \varepsilon/4)$

if *some* $a_i \neq \perp$ **then**

  let $y = x_{max}(s_i)$ and the associated aggregator $S' = s_i$

let $\{a'_k\} = \mathsf{Sparse}(t, \{Q'_k\}, 3\alpha, 1, \varepsilon/4)$

if *some* $a'_i \neq \perp$ **then**

  if $s_i \succ S'$ **then**

    let $y = x_{min}(s_i)$ and the associated aggregator $S' = s_i$

let $\{a''_k\} = \mathsf{Sparse}(\{t, \{Q''_k\}, 3\alpha, 1, \varepsilon/4)$

if *some* $a''_l \neq \perp$ **then**

  for *each* $0 \le j \le n$ **do**

    let strategy profile x^j be defined as

$$x_i^j = \begin{cases} x_{max}(s_l)_i \text{ if } i \le j \\ x_{min}(s_l)_i \text{ otherwise} \end{cases}$$

  let query $Q'''_j = S(x^j)$

  let $\{a'''_j\} = \mathsf{Sparse}(t, \{Q'''_j\}, \alpha + \gamma/2, 1, \varepsilon/4)$

  if *some* $a'''_{j'} \neq \perp$ **then**

    if $s_l \succ S'$ **then**

      let $y = x^{j'}$

  if y *is defined* **then**

    Output y

  else

    Abort

We now prove that Algorithm 6 computes an approximate equilibrium with quality score close to OPT. We know that with probability at least $1-\beta$, all four instantiations of Sparse have error at most

$$err = \frac{100\gamma(\log(2Wn) + \log(8/\beta))}{\varepsilon} \leq \alpha,$$

by our assumption on α. For the rest of the argument, we assume this level of accuracy, which is the case except with probability β.

Suppose that the algorithm outputs a strategy profile $\vec{y}$ from the first two instantiations of Sparse. Let m be the index of the corresponding query, then

$$|S(\vec{y}) - s_m| \leq 3\alpha + err \leq 4\alpha.$$

Since each player is playing a $(2\alpha+\gamma+\zeta)$-aggregative best response to s_m in the profile $\vec{y}$, we know by Lemma 3 that each player is at least playing a $(10\alpha+\gamma+\zeta)$-aggregative best response to $S(\vec{y})$, and so by Lemma 2, a $(10\alpha+2\gamma+\zeta)$-best response.

Suppose that the algorithm does not output anything in the first two instantiations of Sparse. By the accuracy guarantee, we know that for each $s_i \in \{s_1, \ldots, s_{2W/\alpha}\}$,

$$|S_{max}(s_i) - s_i| \geq 2\alpha \text{ and } |S_{min}(s_i) - s_i| \geq 2\alpha.$$

Furthermore, we know there must exist a ζ-approximate pure strategy Nash equilibrium $\vec{x}^*$, so each player in $\vec{x}^*$ must be playing a $(\zeta+\gamma)$-aggregative best response to $S(\vec{x}^*)$ by Lemma 1. Then by Lemma 3, the players are playing a $(2\alpha+\gamma+\zeta)$-aggregative best response for some $s_l \in Z$. For such an s_l, it must be the case that $|S(\vec{x}^*) - s_l| \leq \alpha$ and $\vec{x}^* \in X(s_l)$, so,

$$S_{min}(s_l) < s_l < S_{max}(s_l),$$

otherwise the first two instantiations of Sparse would have output $x_{max}(s_l)$ or $x_{min}(s_l)$. The third instantiation of Sparse would find us such an s_l. Now as with Theorem 6, we can find a strategy profile $\vec{z}$ with the last instantiation of Sparse such that $\vec{z}$ is between $x_{max}(s_l)$ and $x_{min}(s_l)$ and,

$$|S(\vec{z}) - s_l| \leq \alpha + \gamma/2.$$

By Lemma 2 and Lemma 3, each player in $\vec{z}$ is playing a $(4\alpha+2\gamma+\zeta)$-aggregative best response to s_l, and so a $(4\alpha+3\gamma+\zeta)$-best response.

Let $\vec{y}'$ be the equilibrium in $\mathcal{E}(\zeta)$ that gives quality $q(S(\vec{y}')) = \text{OPT}$. Such an optimal strategy profile $\vec{y}'$ would be among the strategy profiles that our algorithm searches. Either we output the profile $\vec{y}'$, or we found a different profile $\vec{x}$ associated with a discretized aggregator of higher quality score (because we enumerate aggregators with higher quality q first). Let s' be the associated aggregator to our output profile $\vec{x}$.

Note that $|S(\vec{y}') - s''| \leq \alpha$ for some discretized aggregator $s'' \in Z$ and also $|S(\vec{x}) - s'| \leq 4\alpha$. Because the order in which we iterates the aggregators gives priority to higher quality, we know that $q(s') \geq q(s'')$, so $q(S(\vec{x})) \geq \text{OPT} - 5\alpha\lambda$. □